ENCODING OF RELATIVE LOCATION OF INTENSITY CHANGES IN HUMAN SPATIAL VISION

Ian R. Paterson

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

1992

Encoding of relative location of intensity changes in human spatial vision

Ian R. Paterson

August 1990

University of St. Andrews
Scotland

Th A1479

Contents

ABSTRACT

The psychophysical experiments and numerical modelling reported in the present study are an investigation into the encoding of relative location of intensity changes in the human visual system. The study attempted, successfully, to explain some geometric illusions resulting from closely spaced image features ('crowding'), and determined the nature of information necessary for making judgments about the separation of intensity changes for different stimulus configurations. Experiments performed fell into two basic categories; those concerned with spatial interference, and studies of spatial interval judgments. The first set of experiments, studying spatial interference with relative localisation for intensity changes, was based on measurements made with stimuli composed of lowpass filtered bars and edges. The most successful model, which accounted for all of the data, was Watt and Morgan's (1984, 1985) MIRAGE; the results suggest that a good explanation of some geometric illusions can be derived using the principles of low-level vision. Spatial interference is strong evidence for combination of information across spatial scales, and the MIRAGE algorithm makes some highly accurate predictions. Relating the separation of image features is a fundamental task for the visual system, but there is no clear understanding of what information the system has available to perform this task. The second set of experiments explored the perception of separation, and precision of judgments of separation, for bars with a variety of orthoaxial contrast profiles. The data indicate that information is combined across spatial scales (as in MIRAGE) under certain circumstances in making separation judgments; this combination of information across scale occurs when the information on the scales combined is in agreement (ie. all scales have some task-related information), but when variance is added on coarser scales which is not relevant to the task, the system is capable of selecting the finest scales of filters available, and using only the information in the finest scale. This adaptive scale-selection process operates even at very brief exposure durations.

Acknowledgements

I am deeply indebted to Dave Heeley, for his patience and friendliness during my time at St. Andrews and beyond; he has an enthusiasm for the study of vision which is quite infectious and he has always been a source of inspiration and good ideas. My thanks go to Prof. David Milner and Prof. Michael Morgan for their time and care in the examination of this thesis, and their suggestions for its improvement. I thank Roger Watt for many recent discussions on the topic of early vision, which have greatly facilitated my understanding of what exactly we can hope to find out by doing psychophysical experiments. Prof. Brian Timney in London, Ontario was a kind and generous surrogate supervisor whilst I was over in Canada. I owe a debt of gratitude to my patient and long-suffering subjects, particularly Lesley Clarke, who suffered more than any other; to her I owe the most, since she struggled to like me and encourage me when I had become a monster locked in a room poring over an unfinished thesis; without her help, I would not have been writing this now.

Copyright Declaration

Declaration for the degree of Ph.D.

I, Ian Robert Paterson, hereby certify that this thesis has been composed by myself, that it is a record of my own work, and that it has not been accepted in partial or complete fulfilment of any other degree or professional qualification.

Signed: Date: 31 Aug 1990

I was admitted to the Faculty of Science of the University of St. Andrews under Ordinance General No. 12 in October, 1986, and as a candidate for the degree of Ph.D. in 1990.

Signed: Date: 31 Aug 1990

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate to the degree of Ph.D.

Signature of Supervisor: Date:

17th June 1991

Chapter 0

Introduction

0.1: Preamble

One view of how we should study vision holds that only analysing what information is available to the visual system to perform a task is liable to deliver meaningful results (Watt, 1990; Marr, 1982). However, there is a venerable tradition of trying to understand something about our psychological experience of the world in terms of theories of one sort or another, and it may be that our understanding of vision can be increased by studying some of the subjective aspects of vision, at least if done carefully enough, with due regard for the perversity of human psychology.

The study is a psychophysical investigation of the perception of relative spatial location. A great deal is known about how the human visual system assesses the relative location of image features, but there are some aspects of relative localisation which have not yet been fully understood, and to date no comprehensive explanations exist of the geometric illusions. This study was an attempt explain some very simple geometric distortions using soundly based principles of low-level vision. The experiments reported here involved measuring location biases and precision on relative localisation tasks. From spatial interference experiments, it was found that changes in perceived location induced by the presence of flanking features can be explained by MIRAGE (Watt & Morgan, 1985); interference with both vernier acuity and direction of displacement discrimination, the difference between the data from these tasks, the differing pattern of spatial interactions between bars and edges, and the effects of blur on these tasks are explained by MIRAGE in elegant fashion.

A series of spatial interval experiments were performed which involved systematic or random manipulation of the degree of asymmetry of the luminance profile of bars delimiting a spatial interval. The results of these experiments eliminated features in the stimulus intensity distribution (centroids, inflexions, threshold edges and luminance maxima) as delimiters of perceptual intervals. MIRAGE explains the phenomenology of the results of several of these experiments, but it cannot be reconciled with data from separation discrimination tasks where the information present on the coarser spatial scales is varied randomly in each stimulus presentation. Separation discrimination is apparently a counter-example to 'coarse-to-fine' processing in the visual system, with information from the finest scales available independently of information on coarser scales even with very brief exposure durations, if the task demands it. A dissociation is found between the phenomenology of perceived separation of luminance features and the information necessary to judge their separation. Combination of information across spatial scales was implied by the pattern of biases measured for stimuli without spatial perturbation, but could not be reconciled with the precision of separation judgments for stimuli with spatial perturbations. When the stimuli does not vary randomly in its spatial configuration, the system appears to combine information across several scales, but when there is externally applied noise on the coarser scales that the system can identify as noise, then it can restrict itself to using information present only on a fine scale. This closely parallels an observation by Burbeck & Yap (1990), indicating that the visual system can choose to accept information only from spatial scales which deliver a reliable signal which is task-related, indicating a role for selective attention even at the earliest stages of edge localisation.

0.2: Early vision

The general framework of this investigation is based upon ideas developed by David Marr (1976, Marr & Hildreth, 1980; Marr, 1982) and Watt and Morgan (1984, 1985) who applied his approach to the study of early vision. Marr (1982) stressed that a complex system may only be understood if a number of levels of analysis are simultaneously brought to bear on the system. He argued that for a problem as complex as biological vision to be understood, it necessarily should be understood at the levels of computational theory, algorithms and hardware. His main point was that the computational theory should guide research at the other levels. The computational theory states what is being computed, and why it is being computed. It should specify what sort of representations are necessary (ie. what information is made explicit at a given stage in processing) and how the information in a representation is manipulated by the processes which act upon it. The algorithmic level gives details of how to implement the theory, which is likely to be hardware-dependent; the computational theory ought to be hardware-independent. Marr believed that given the complexity of the neural architecture, an understanding of what the brain did would not arise by direct examination of what it looked like and what it appeared to be doing, but rather by a rigorous study of the *information necessary for a task* and how the task might be carried out. Central to Marr's ideas is the belief that early processes are largely data-driven; the visual system is thought to be concerned with squeezing as much information necessary to make a correct interpretation out of the image before any higher-level processes are involved. A representation is simply a means of encoding information- a description of some sort, which may make explicit properties of the visual scene which are not obvious upon examination of the intensity distribution of the image (eg. the location of texture boundaries, for instance).

Marr regarded early vision as a set of computational processes which detect and represent intensity changes in the image, extract information about local geometric structures in the image, and detect and analyse the effects of illumination (eg. determining whether a surface is opaque or transparent). He viewed the ultimate aim of early vision as construction of a visible surfaces representation (the so-called '2.5 D sketch', which is a viewer centred description of the visual image), in which surface attributes such as orientation and relative distance from the viewer might be made explicit, based upon combining information from stereo, motion, shading, shadows, texture, intensity and chromatic changes. The most primitive representation in Marr's scheme is the primal sketch (which Marr subdivided into 'raw' and 'full' primal sketches), which is retinotopically organised. The primal sketch is composed of a set of spatial primitives. In Marr's scheme, spatial primitives, are cousins to Lotze's (1884) retinal mean local signs. Spatial primitives encode attributes of image features as edge/ boundary segments, bars, feature terminations and 'blobs' (doubly terminated bars). These primitives in turn are acted upon by early grouping processes (which embody prior knowledge that the visual system has about the geometry of the visual world). Marr and Hildreth (1980) devised a scheme for detecting and representing intensity changes in the retinal imge, based on a number of bandpass spatial frequency filters which are used to analyse the retinal image in parallel. The outputs of the filters are combined and features in the combined output extracted to give a 'raw primal sketch' of the image. Watt and Morgan (1984, 1985) applied a related model to data from extensive psychophysical experiments to explain performance on relative localisation tasks. These models are described in the next chapter. Most models in computational vision which have as initial stages in their operation filtering by spatial frequency bandpass filters; extensive work on multiple spatial frequency channels and spatial contrast detection, started by Campbell and Robson (1968), indicates the existence of multiple spatial frequency bandpass filters at each retinal locus in the human visual system. The evidence for this notion is described in the next chapter.

Chapter One

Studying early vision

1.1: Overview

The purpose of this chapter is to discuss psychophysical and theoretical work relating to how the human visual system detects and represents elementary image features, and computes spatial relations amongst such features. Psychophysical data and theory relating to detection and discrimination of contrast-modulated patterns is discussed in the sections 1.2 through 1.6. The literature is vast, even when attention is confined to static, achromatic luminance modulated patterns, and the account here is inevitably highly selective. Section 1.2 presents empirical evidence for the existence of multiple spatial frequency bandpass filters in early vision, derived from studies of the visual system's sensitivity to spatial modulations of luminance; the existence of such bandpass filters is central to most theories of low-level vision.

There are many different pattern acuities which have been investigated- a pattern acuity is simply the limit on *precision* in judgments about elements of the geometry of one or more pattern(s). Measures of precision can be highly informative, since precision (the reciprocal of the subject's error variance) is a direct measure of the amount of information the system has available in performing the task. Ideally, we wish some mathematical model, which should explain why the human observer shows the pattern of information loss which they do (Watt & Morgan, 1983b). A major part of the task of psychophysics is to discover what cues subserve psychophysical judgments, a cue simply being an information source. Section 1.3 describes results from studies of vernier acuity and orientation discrimination. The following section, 1.4, is an account of data from spatial interval and spatial frequency discrimination experiments; these sorts of discriminations are relative localisation tasks which display fundamental differences to vernier acuity. Unlike vernier acuity, spatial interval discrimination is remarkably robust to changes in the spatial configuration of the target. Psychophysical evidence for the existence of spatial primitives in human vision is discussed in section 1.5. Section 1.6 samples some theoretical accounts of early visual processing; section 1.7 is a very brief account of some of the neurophysiological research which is relevant to psychophysical studies of early vision. Whilst several workers have attempted to make close links between physiological and psychophysical results, it seems ambitious at this stage in our knowledge to do so. To date, there are no neurobiological linking hypotheses which allow reconciliation of psychophysical and neurophysiological results. On the other hand, physiology and anatomy can identify constraints on what sorts of representations and processing occur in biological visual sytems. Additionally, neurophysiological results have been used to generate models of relative localisation of image features, and some of these models are briefly examined in this section.

1.2: Spatial contrast detection and multiple channels theory

This section discusses briefly some of the psychophysical evidence in favour of multiple, spatially coincident and relatively independent spatial filters in human vision. Most empirical evidence on the nature of 'channels' has come from studies of spatial contrast detection (using adaptation, subthreshold summation, masking and other related paradigms) where the contrast of the lowest contrast pattern that can just be seen is determined as a function of the stimulus parameters. The most difficult problem with such studies is that it is hard to generalise the results to suprathreshold contrast levels. There has has been an extensive and systematic study of the properties of early spatial filters or 'channels'. The major analytic tool has been linear systems theory, which has had a considerable influence upon thinking about spatial vision.

1.2.1: The multiple spatial frequency channel model

There is ample evidence that neural mechanisms exist in human vision, which correspond to bandpass spatial frequency selective mechanisms (Campbell & Robson, 1968; Blakemore & Campbell, 1969) and which are also spatially localised (Wilson & Bergen, 1979). Several different spatial scales are believed to co-exist at each retinal location (Graham, Robson &Nachmias, 1978; Wilson & Bergen, 1979).

The proposal that the visual system might be transmitting independent spatial frequency signatures of the retinal image was made by Campbell and Robson (1968). They determined how well subjects could distinguish a square wave modulation of contrast from a sine wave grating; they concluded that "a square wave grating is perceived to be different from a sine-wave grating when the third harmonic of the square wave independently reaches its own threshold." They explicitly compared the visual system to a linear system and proposed that early vision literally performs a Fourier analysis of the retinal image. Blakemore and Campbell (1969) determined contrast thresholds for gratings following adaptation to a high contrast grating, and found that thresholds were elevated only for frequencies near to the adapting frequency, which they interpreted as a selective impairment of a subset of the visual system's spatial frequency channels. Graham and Nachmias (1971) demonstrated that the detectability of a compound grating composed of more than one sinusoidal component did not depend on the relative phase of the components. Sachs, Nachmias and Robson (1971) interpreted their data as evidence that unless gratings are very similar in frequency, they are detected by independent spatial frequency selective filters.

1.2.2: A space-variant single-filter model

All of the data presented above could be explained by a model in which the frequency components of spatially extended patterns like gratings are independently detected at different visual field locations. From psychophysical (Green, 1970) and physiological (eg. Hubel & Wiesel; 1962, 1974a, 1974b) evidence, the visual field is known to be markedly inhomogeneous, with a general increase in the spatial scale of receptive fields at more peripheral locations. It is possible that the relatively independent detection of widely separated frequency components requires the components to be detected at different retinal locations. The sensitivity of central vision to fine spatial detail (ie. information conveyed by high spatial frequencies) has long been known, and it is possible that low spatial frequency components could be detected at more peripheral visual field locations. A model incorporating the inhomogeneity of the visual field was devised by Limb and Rubinstein (1977). This model incorporated a single linear filter at each retinal location, but the scale of the filter increased with eccentricity. Since the filter was linear, its response could be described by a convolution with the input pattern. Their filter parameters were derived from measurements of the sensitivity of the visual system to two line patterns separated by different distances, which they used to determine a one-dimensional line-spread function (spatial weighting or sensitivity function) at six different eccentricities (0-8.3 deg arc). The filter profile had a large number of free parameters to ensure a good fit to the sensitivity data. They then predicted thresholds for other classes of patterns (eg. bright/ dark bars, gratings) using this space-variant single filter model; the model was quite successful in explaining all the data from their experiment. Energy summation across spatial frequency is predicted by a single lienar filter model; presenting stimuli with substantial energy at two widely separated frequencies can be used to test whether energy is combined across spatial scales.

1.2.3: Evidence for multiple spatial filters at one locus

Subsequent studies showed that there is more than one scale of spatial filter at each retinal location. Graham, Robson and Nachmias (1978) measured summation of sinewave gratings of different frequencies, for patches of grating windowed with a Gaussian to localise the stimulus in a particular retinal region (presented either foveally or 7.5 deg peripherally), in order to minimise the confounding effects of retinal inhomogeneity on detection. The grating patches were composed of gratings of 2 and 6 c/deg, in either 0 or 180 deg relative phase (peaks-add and peaks-subtract, respectively). The detectability of the f-3f compound is less than would be expected if one linear filter were present at each retinal locus (as in the model of Limb and Rubinstein); detection of the different components was slightly better than predicted by a wholly-independent filter model, which they attributed to probability summation across spatial frequency detectors (eg. a filter with a peak frequency of 4 c/deg has some probability of responding to a grating of 6 c/deg or 2 c/deg).

If patterns are presented at contrast threshold, then discrimination may require that the patterns are detected by different filters; ie. that somehow the filters are 'labelled', although this is a strong assumption. Watson and Robson (1981) used a detection/ identification paradigm to examine spatial frequency discrimination for Gabor patches (sinewave windowed with a Gaussian) presented at contrasts near to detection threshold. In this paradigm, the subject is required to indicate which of two temporal intervals contained the grating patch, and whether it was of the higher or lower of the two spatial frequencies being presented within any one session. To give 'perfect' discrimination at contrast threshold (being when the grating is correctly identified every time it is detected), a difference in spatial frequency of about an octave between the patches was necessary. Watson and Robson concluded that about seven classes of independent local spatial frequency detectors were necessary to explain performance on these tasks, if errorless discrimination depended on signals from labelled filters with different peak spatial frequencies.

Wilson and Bergen (1979) developed a highly influential model of threshold spatial vision, based on the operation of four independent spatial frequency filters at each retinal location; this is described in section 1.6. The peak frequencies of the filters in Wilson and Bergen's model ranged from 1-8 c/deg. An experiment by Watson (1982) suggested the need for a wider frequency range of filters than those in Wilson and Bergen's model. Watson presented compound grating patches spatially windowed with a Gaussian; the higher of the two component frequencies in the patch was varied, and the detectability of the compound patch measured. The lower frequency was either 1 or 16 c/deg. Watson measured less summation than would be predicted by the filters from Wilson and Bergen's model, indicating relatively independent detection of 1c/deg and 2.8 c/deg patches was occurring (as well as for 16 c/deg and 32 c/deg patches).

1.2.4 Fourier analysis and early vision

The view of early vision as equivalent to a Fourier analysis of the retinal image was the dominant paradigm in spatial vision for many years. From a computational perspective, this is not really a theory of vision at all. Fourier analysis allows an arbitrary waveform, as long as it is limited by physical possibility, to be re-expressed as the sum of periodic components. Fourier methods transform a waveform expressed as a function of space and/or time into an equivalent representation in the frequency domain. Fourier analysis can describe the response of all systems that are linear and stationary (time-invariant). The Fourier transform of a waveform can be expressed as a sum of sinusoids with different frequencies; each frequency has an associated amplitude (or energy) and phase.

Descriptions in the frequency domain and space domain are formally equivalent; no information is sacrificed by a Fourier transform, but the only useful property of a simple Fourier transform for an image analysis sytem is that the transform gives a shift-invariant representation if (and only if) the phase spectrum is ignored. Since the phase spectrum carries information about local image geometry, this neglect leads to problems. In terms of Marr's approach, there appears to be little made usefully explicit in transforming the image into a frequency representation. It might be argued that with high spatial frequencies generally carrying texture, low spatial frequencies 'shape' there is some useful partitioning of image information; but we do not need an exclusively frequency domain representation to do this- we can do this equally well with multiple spatial scales of representation. One major problem with the work on spatial contrast detection was that it was driven by the idea that early vision performed a spectral analysis of the input, and that the system was essentially linear. The models developed could not be made to generalise from the response of the visual system to threshold contrast patterns to make accurate predictions about visual responses made at suprathreshold contrast levels (Olzak & Thomas, 1987).

1.3: Relative localisation

There are many sorts of acuities, the most frequently encountered being resolution acuity - ability to see fine spatial detail; for instance, to tell that two dots are separate, or see a high spatial frequency grating pattern. The lowest values obtained for resolution acuity are in the order of 30 sec arc, or close to the diameter of a foveal cone. There are other sorts of pattern acuities which yield finer thresholds than this; vernier thresholds, for instance, can be as fine as 3 sec arc (Westheimer & McKee, 1977a); since this is about one-tenth of the diameter of the smallest photoreceptors, the processing underlying vernier acuity is a form of neural interpolation (Barlow, 1981; Morgan & Watt, 1982).

1.3.1: Vernier acuity and mean retinal sign

One familiar result in vernier acuity is improvement in sensitivity to vernier offset as the length of the lines in the target is increased. For abutting targets, there is a limited dependency of vernier acuity on line (or edge) length (Anderson & Weymouth, 1923; Andrews, Butcher & Buckley, 1973; Westheimer & McKee, 1977a). The spatial scale of the target determines the retinal distance over which increasing the feature length can improve vernier acuity. Watt and Morgan (1983b) demonstrated that the region over which target length can improve performance is proportional to the blur of the target features. For instance, with edges having a space constant of 10 min arc, they found that an edge length of 1 deg is necessary for asymptotic performance. For fine line targets, vernier threshold is almost independent of line-length beyond an overall target length of about 10 min arc (Westheimer & McKee, 1977a).

The dependence of vernier acuity on line-length was used by Hering (1899) to explain the precision of vernier acuity, which he saw as resulting from averaging across the length of the lines in a vernier target to derive a retinal mean local sign for each line. This notion of a local sign is based on information in the retinal intensity distribution itself, but more recent forms of this idea are based on information in blurred and differentiated versions of the retinal image, although the fundamental concept is the same- representing an image feature by a 'token' or symbol of some sort, which itself possesses only a limited amount of information. Evidence for spatial primitives comes from studies showing that primitive features of a given sort result in a comparable information loss to that shown by the human visual system in making a judgment- that is, by demonstrating that precision in a given task declines in a manner which would be predicted by supposing that only the information encoded by these spatial primitives is available in performing the task.

1.3.2: Two mechanisms in vernier acuity

Ludvigh (1953) showed that two dot vernier acuity can be as precise as two line vernier acuity, which refutes the notion of a mean local retinal sign as the only explanation of the precision of vernier alignment. It is also possible to use a relative orientation cue in performing a vernier task. Relative to an aligned target, a vernier target with an offset appears as rotated slightly (eg. rotated clockwise if the lower bar is offset to the left of the upper bar). Vernier acuity may reflect the operation of at least two localisation mechanisms, one concerned with extracting an overall orientation for an image feature, and the other with extracting information about the relative location of the components in the target, in the dimension perpendicular to the main axis of the target. Evidence for the existence of two distinct mechanisms involved in vernier acuity was provided by Watt, Morgan and Ward (1983a). In a conventional vernier task (eg. judging whether one line above the other is offset to the left or right), the subject may use both the relative orientation of the target (eg. relative to vertical) and the relative location of the components of the target in judging the direction of vernier offset. If the orientation of the target is varied randomly from trial to trial, then the relative orientation cue becomes unreliable, since a change in perceived orientation is not related to the direction of offset. Performance on the vernier task should then be determined by the operation of the mechanisms which compare local signs for the target elements. Ideally, we would like to show a relation between physical cue amplitude and precision of localisation- ie. to discover what physical cue controls localisation precision; if we degrade a given cue and localisation performance deteriorates in a systematic fashion, we can assert that information carried by this cue is necessary for localisation (since precision of localisation is inversely related to the amount of information used by the system). For a conventional vernier judgment, comparison of local signs will have to occur along the direction perpendicular to the main axis of the target; ie. in the "orthoaxial" direction (Watt & Andrews, 1982). We can degrade the orthoaxial contrast information available to the system by blurring the target. This was performed by Watt, Morgan and Ward (1983a); they blurred the corners of an abutting vernier target, and determined vernier thresholds for various degrees of corner blur. When sharp corner detail is present, randomising the reference orientation (with the overall target length constant at 30 min arc) has no effect on vernier thresholds. Blurring the corner detail leads to a difference between the random and fixed reference orientation conditions, for greater degrees of corner blur. The decline in performance seen in the random slope condition (when no relative orientation cue is present) is consistent with the degradation in the amount of orthoaxial contrast information present in the target. It is therefore reasonable to conclude that this information is used by at least one of the mechanisms involved in vernier acuity.

1.3.3: Orientation discrimination

Orientation discrimination is another sort of pattern acuity which is capable of yielding thresholds in the 'hyperacuity' range; a change in orientation for a pattern of as small as 0.5 degrees rotation can be detected by the visual system, which for short lines yields thresholds of less than a cone-diameter, ie. in the hyperacuity range. Despite a superficial similarity to vernier judgments, there are some quite striking differences between orientation discrimination and vernier acuity.

Like vernier acuity, orientation discrimination is dependent on line length (Andrews, 1967) with improvements reported at line lengths of at least one degree of arc for an abrupt 1 msec ("transient") presentation, and up to 30 min arc for "continuous" presentation (unlimited inspection time). Spatial interference with orientation discrimination has a similar pattern to interference with vernier acuity (Westheimer, Shimamura & McKee, 1976), with optimal interference occuring when two symmetrically disposed flanking lines are situated some 3-6 min arc distant from the target line.

Orientation discrimination improves over part of the contrast range, but performance saturates at relatively low contrasts of about 5-10% (Regan & Beverley, 1985; Skottun, Bradley, Sclar, Ohzawa & Freeman, 1987). This makes it somewhat different to vernier acuity, which is dependent on contrast for the whole of the contrast range (Watt & Morgan, 1984). Another difference is that orientation discrimination thresholds are largely independent of spatial frequency (Heeley & Timney, 1987), unlike vernier acuity, which is related to the spatial scale of the targets to be localised (Watt & Morgan, 1984; Bradley & Skottun, 1987). There is some evidence that orientation discrimination is unaffected by the presence of a large spatial frequency difference between the stimuli to be compared (Bradley & Skottun, 1984), although this may be artefactual (Heeley, pers. comm.). Both orientation discrimination and vernier acuity display meridional anisotropies (McKee & Westheimer, 1978; Heeley & Timney, 1987; Orban, Vandenbussche & Vogels, 1984; Corwin, Moskowitz-Cook & Green, 1977), which may reflect the use of a relative orientation cue in the vernier task becoming less effective at oblique orientations.

1.4: Spatial frequency and spatial interval discrimination

One school of thought (Olzak & Thomas, 1987) regards spatial frequency discrimination as a degenerate case of size discrimination (in which the stimulus is highly redundant), since the width of bars in a grating are inversely related to the spatial frequency. This section treats spatial frequency discrimination as a spatial interval discrimination between bars of the same contrast polarity in the grating target. Spatial period discrimination reaches an asymptote at 1.5 cycles of a sinewave grating (Hirsch & Hylton, 1982), suggesting that optimal performance can be based on just one spatial interval (eg. between two bright bars in the grating pattern), although alternative strategies may be triggered by using stimuli other than simple sinewave gratings in the discrimination task; the system can use information contained in more than one spatial interval if the task demands it. Heeley (1987) determined spatial period discrimination performance for grating patterns with random frequency modulation; performance is better than would be expected if the system simply randomly selected one spatial interval (eg. between two luminance maxima), and compared it with another similar interval randomly chosen from the comparison pattern. Rather, the system is capable of integrating a series of measurements of spatial period (ie. a spatial averaging process) to deliver an estimate of spatial period which is more precise than if the judgment were based on one randomly chosen interval. The system probably deals with these frequency-modulated patterns as if they were a form of linear texture, and the judgment is akin to estimating mean 'blob' width in the pattern.

1.4.1: Precision and perturbation

The precision of both spatial frequency and interval discrimination is surprisingly resistant to changes in the luminance profile of the retinal image that are known to be important in detection tasks, which immediately suggests difficulties in attempting to explain precision of such discriminations in terms of the properties of low-level detection mechanisms (eg. Wilson & Gelb, 1984). Fractional thresholds for spatial frequency discrimination are relatively constant over a wide range of spatial frequencies (Campbell, Nachmias & Jukes, 1970), in contrast with the dependence of sinusoidal grating visibility on spatial frequency (Blakemore & Campbell, 1969). Fractional difference thresholds are similar across a wide range of spatial frequencies at very low contrasts if the stimuli are equated for visibility (Thomas, Gille & Barker, 1982; Watson & Robson, 1981). Campbell et al (1970) found no significant difference in discrimination thresholds for square versus sinewave gratings, and no superiority for binocular over monocular presentation. Spatial frequency discrimination thresholds are known to be independent of contrast over most of the contrast range. Caelli, Brettel, Rentschler and Hilz (1983) measured frequency discrimination thresholds at contrasts ranging from 5 to 25 % contrast, at different reference orientations and frequencies, and found no effect of contrast, replicated by Skottun, Bradley, Sclar, Ohzawa and Freeman (1987) for a wider contrast range. Morgan and Regan (1987) demonstrated that precision of spatial interval discrimination is unaffected by a contrast randomisation procedure (for separations greater than 2.5 min arc). Similarly, edge blur width discrimination, wherein subjects are required to estimate the spatial extent of a blurred edge, is unaffected by contrast randomisation (Hess, Pointer & Watt, 1989).

Spatial interval discrimination is relatively robust to changes in the spatial configuration of the target. Two point separation discrimination is as good as that for line separation (Westheimer & McKee, 1977a). Morgan and Ward (1985) determined spatial interval thresholds for intervals with flanking features placed a randomly varying distance either side of the interval delimited by two lines; this manipulation did not affect the precision of interval judgments.

Spatial interval discrimination is also apparently insensitive to the spatial frequency composition of the individual features delimiting the interval. Burbeck (1988) devised spatial interval targets which were delimited by either lowpass Gaussian bars, or high-frequency bandpass modulated bars, or one of each. She reported that interval discrimination was as accurate with the mixed frequency pairs as with those matched in frequency. Likewise, Toet and Koenderink (1988) found that the frequency content of a Gabor patch had no effect upon the thresholds for spatial bisection and displacement tasks with three "blob" targets, whose intensity profile was a two-dimensional Gaussian function, presented at threshold contrast. Bradley and Skottun (1984) and Burbeck and Regan (1983) both reported an apparent insensitivity of spatial frequency discrimination to the presence of a 90 deg. orientation difference between the gratings to be discriminated. Additionally, spatial period discrimination is largely unaffected by duration of the inter-stimulus interval; Regan (1985) showed no decline in precision of period estimation for inter-stimulus intervals of up to twenty seconds. Judgments are robust to variations in the relative distances between the comparison stimuli (Campbell, Nachmias & Jukes, 1970), and contrast randomisation (Heeley, in press). Spatial interval discrimination can be as precise in a single interval judgment, where the subjects must compare the width of the target stimulus with some internal standard, acquired through feedback (Westheimer & McKee, 1977a).

Beyond small separations, accuracy of interval discrimination is not affected by the polarity of contrast of the features delimiting the interval. Discrimination of small and large separations may be performed by somewhat different processes; Levi and Westheimer (1987) showed that interval discrimination for lines of different polarities is poorer than for lines of the same polarity when the separation is small (ie. less than about 4 min arc), and this is supported indirectly by the result of Morgan and Regan that contrast randomisation of the bars in the interval target did affect interval discrimination at the smallest separation they used (2.5 min arc).

1.5: Psychophysical evidence for spatial primitives

To demonstrate that the visual system uses spatial primitives, it is necessary to show that extracting and representing certain sorts of features result in a comparable information loss to that suffered by the human visual system (Watt & Morgan, 1983b). A psychophysical result which provided circumstantial evidence favouring extraction and representation of spatial primitives in early vision was that of Westheimer and McKee (1977a); they demonstrated that subjects could apparently estimate the location of the centroid the intensity profile of a narrow bright bar in a vernier task, with a precision which was equivalent to that for a conventional vernier judgment with fine, bright lines. They presented narrow bars (less than 3 min arc wide); one of the two bars had a luminance asymmetry, and subjects were simply asked to perform a vernier alignment of the two bars. Subjects could accurately discriminate the relative location of the centroids of the two bars (with a precision equivalent to 'normal' vernier acuity for symmetric bars) without being aware that the bars were asymmetric in luminance. Whether the subjects were actually estimating the location of the centroid of the retinal intensity distribution was not clear- it is possible that some other feature (which happened to fall close to the location of the centroid of the intensity distribution for these stimuli) was being used, but they did not distinguish between any alternative candidates.

1.5.1: Zero-crossings

The first empirical demonstration of the perceptual significance of zero-crossings in the second spatial derivative of a smoothed version of the image was provided by Watt and Morgan (1983a). They measured the perceived width of narrow, bright bars with a luminance asymmetry, relative to a symmetric bar. Luminance asymmetries were created by having two very narrow bars with differing luminances spaced below the resolution limit. Subjects could discriminate the 'paired' bar from a single bar, when the change in the separation of the peaks of the two bars was much less than the theoretical resolution limit for these two stimuli (about 1 min arc); subjects were discriminating the width of the paired bar and the single bar, and these width thresholds increased as the degree of luminance asymmetry in the 'paired' bar increased.

They examined whether zero-crossings in the second derivative of a smoothed retinal image could explain the thresholds for changes in perceived width of these asymmetric bars. They simulated the neural response with a DOG filter (centre-surround ratio 1:1.75, excitatory centre space constant 22 sec arc). Zero-crossings accounted for the width thresholds obtained with these asymmetric stimuli. Other candidates for edge primitives they discounted were threshold edges (places in the image where the luminance exceeds some threshold value) and luminance maxima; these intensity features made predictions which were inconsistent with the results of vernier acuity for targets composed of one symmetric and one asymmetric bar, relative location in the vernier task depending on either the mean of the intensity distribution, or the zero-crossings in the 2nd derivative of the smoothed image. They could not choose between the mean/ zero-crossings models on the basis of the vernier acuity data, since the largest luminance asymmetries they could create did not lead to predictions discriminable within experimental error, but only the zero-crossings model could explain the bar-width discrimination data.

1.5.2: Centroids and extrema

In a further set of experiments Watt and Morgan (1983b) examined the suggestion of Marr and Hildreth (1980) that edge blur and contrast could be computed from the scale of filter responding to an edge, and the gradient of the filter response at the zero-crossing (in the direction perpendicular to the 'edge' alignment). Marr and Hildreth's edge primitives were zero-crossing segments and associated gradients of the normals to these segments, computed by filtering the image with several scales of second-order filters. Watt and Morgan determined thresholds for detecting a change in the spatial extent of blurred step edges, which had either Gaussian, rectangular or cosinusoidal blurring functions. They noted that sensitivity to an increment of blur added to a sharp edge is less than an increment added to a slightly more blurred edge, which was one of the critical pieces of evidence for their theory of spatial primitives MIRAGE (Watt and Morgan, 1984, 1985); this indicates the addition of some degree of internal 'blur' to an edge, which is consistent with the responses of larger scale filters being combined in some way with the smaller scale filters to increase the effective neural blur. For blur extents of less than about 3 min arc space constant, threshold increases with decreasing blur extent.

Firstly, they determined blur extent thresholds as a function of contrast, for contrasts from 10-80%, for a Gaussian blurred edge with a space constant of 2.5 min arc. Thresholds improved with contrast, and gave a inverse square root relationship between contrast and blur extent thresholds, ie. a power law with an exponent of -0.5. They also measured blur extent thresholds for the Gaussian, rectangular or cosinusoidal blur types as a function of the pedestal blur. They obtained u-shaped functions of threshold versus pedestal blur; the rising portion of the functions for the three types of blur had a power law exponent of 1.5 (1.0 corresponds to a Weber's law-type relationship). This is an unusually large exponent for a spatial acuity, since spatial frequency discrimination (for instance) gives Weber's law-like behaviour. If the blur extent is defined as the separation of the peak and trough in the second derivative of the retinal stimulus, and the threshold expressed as the change in separation of these features which can be distinguished at threshold, then the three functions from the different blur functions coincide, which

Watt and Morgan note is a remarkable coincidence. For separations of the extrema in the second derivative of the retinal stimulus of greater than 5 min arc, a power law with an exponent of 1.5 fits the data for the three different blur types and the functions have the same height above the x-axis. The simplest way of describing their data is the following function,

$$S_b = \frac{A \,.\, S^{1.5}}{\sqrt{C}} \quad \text{for values of } S > 5 \text{ min arc}$$

where S_b = the blur extent difference threshold,
A = a constant
C = contrast
S = pedestal blur

They dismissed Marr and Hildreth's suggestion that edge blur could be computed from the scale of filter responding, since this scheme would predict that edge blur thresholds for rectangular blur of medium and greater blurs would be much poorer than they were (since the smaller scale filters would not give zero-crossings in response to moderately blurred rectangular edges). The gradient at the zero-crossing in the second derivative of the retinal stimulus cannot be a cue to blur extent; since this is proportional to blur extent and inversely proportional to contrast, then the power law exponents for the effects of pedestal blur and contrast should be equal in magnitude but opposite in sign (whereas they are different in reality by a factor of three).

They also measured localisation thresholds for blurred edges, using a vernier task. Vernier thresholds were measured for edges with Gaussian, rectangular and cosinusoidal blur, as functions of both contrast and edge blur. Vernier thresholds showed an inverse square root dependency on contrast, and a square root relationship with edge blur. To explain this data, a very simple statistical model was devised; they equate the task of finding the extrema of a zero-bounded filter response distribtion, which is a portion of the blurred second spatial derivative of the retinal image, with finding its first moment (centroid). By analogy with the process of estimating the mean of a distribution, where the precision of the estimate of the mean is given by the standard error of the mean (the standard deviation divided by the square root of the number of samples), the precision with which this centroid can be located will depend on both the dispersion of this distribution and on the area of this response distribution, the dispersion of the distribution being equivalent to the standard deviation, and the area under the distribution the number of samples (from this analogy). The area of the distribution will depend on both contrast and edge blur, but they assume the width depends only on edge blur.

Their working is shown below.

From their analogy,

$$S = k.\frac{W}{\sqrt{A}}$$

where S is the theoretical localisation threshold (the s.d. of the error response distribution)

k = a constant

W = the width of the zero-bounded response distribution

A = the area of the zero-bounded response distribution

Since the area depends on both contrast and width,

$$A = k'.W.C$$

we can substitute in the above equation for A and simplify, giving

$$S = c.\frac{\sqrt{W}}{\sqrt{C}} \text{ where c is a constant}$$

which was the observed relationship between the stimulus parameters and the psychophysical threshold; expressing localisation thresholds as a function of stationary point width (at 80% contrast) gave a threshold which showed a square root dependency on the stationary point width. The inverse square root relationship of threshold to contrast has been discussed above.

1.5.3: Interpreting primitives

They also discussed how the visual system might actually use primitives such as the centroid of a region of filtered response to represent features such as bars and edges. For isolated bars or edges, there is not generally a computational problem in deciding which positive centroid goes with which negative centroid, but this will not be generally true for natural images. They propose that a region of zero-filter response between pairs of centroids of opposite sense is necessary to allow segregation of the features corresponding to distinct edges or bars; in other words, to solve the correspondence problem for the features which together represent the edge. This is a neat way out of what could be a tricky computational problem, and is supported indirectly by several brightness illusions, which can be explained by supposing that the system fails to register the "true" nature of the visual image if there is either an absence of a zero-region between positive and negative centroids which do not belong to the same edge, or if there is a fortuitous region of zero-response between features which should belong together. In the Chevreuil illusion, consisting of two closely spaced step edges, three edges are seen instead of two, but this is only true for small spacings of the two step edges; consistent with the system failing to resolve a region of zero-response between two centroids of opposite sense which do not belong to the same edge. Mach bands can be seen for ramp edges, but not at very small sizes of ramp, consistent with the system adding in extra edges of opposite polarity in situations where features which should be kept separate are incorrectly combined into edge-primitives.

1.5.4: Perceived size and spatial extent

Many studies have looked at changes in perceived spatial frequency; for instance, following adaptation (Blakemore, Nachmias & Sutton, 1971; Heeley, 1979) or shifts in perceived frequency due to changes in contrast (Georgeson, 1980). The rationale for these experiments has typically been to demonstrate that perceived spatial frequency or size can be explained by the relative activity of spatial frequency tuned channels. Gelb and Wilson (1983a,b) tried to explain perceived size in terms of such frequency channel models, but found them unable to account for many observations relating to the perception of the size of spatially localised patterns.

Gelb and Wilson (1983a,b) studied changes in the perceived size of spatially localised Difference-of-Gaussian (DOG) patterns produced by varying contrast, temporal modulation or by the presence of a masking sinewave grating. The effect of reducing the contrast of a DOG pattern is to reduce its perceived size relative to a high contrast standard, also observed by Georgeson (1980) for Gaussian bars. These results held under both sustained and transient temporal modulations. To attempt to explain their results, they employed a model for perceived size based on earlier formulations by Klein, Stromeyer and Ganz (1974) and Georgeson (1980), using filter space constants and sensitivity parameters from Wilson and Bergen (1979).

The size-index measure they computed is:
perceived size = $\sum w_i \cdot r_i$
where $r_i = R_i / \sum R_i$
R_i represents the response of the ith filter in Wilson and Bergen's model,
r_i is the proportional contribution of the ith filter to the total response,
w_i is a weighting coefficient which is proportional to the space constant of the ith filter.

By weighting the filter responses in the above fashion, Gelb and Wilson are incorporating Watson and Robson's (1981) labelled detectors idea, which asserted that perceived size depended on the 'identity' of the filter detecting a stimulus, although Watson and Robson were studying discrimination of Gabor patches at threshold contrast, and the idea could only be used at suprathreshold levels with additional assumptions about how filter outputs are combined in determining perceived size. The simplest assumption is linear addition of the filter outputs, which is expressed in the above summation to compute a size index.

Using the parameters for filter space constants from Wilson and Bergen (1979) they did not obtain a good fit to the data. Of course, the above size-index uses only information in the frequency amplitude spectrum; it computes a sum of filter responses for filters centred under the stimulus, and takes into account only the amplitude of the filter outputs. Gelb and Wilson suggested that the distribution of filter responses across space might also be important in determining perceived size.

Their second study looked at the effect of masking by sinewave gratings on the perceived size of a DOG pattern. The perceived width of an unmasked DOG standard was compared with a test pattern superimposed on a sinewave of variable spatial frequency.

Perceived size was dependent on the phase of the grating mask relative to the centre of the DOG pattern. Gelb and Wilson note that a zero-crossings model, eg. based on Marr and Hildreth's (1980) scheme, can explain their data (for the condition where mask and target have the same orientation), whereas the size-index model (based on the energy in local spatial frequency detectors) fails to account for the data. Perceived size of these DOG patterns was also influenced by an oblique mask (which randomised the phase relationships between the test pattern and the mask), and this effect was spatial frequency dependent, which Gelb and Wilson argued represented evidence against a spatial primitives type approach- however, it is not clear that this result could not simply be due to an illusion of simultaneous size contrast (itself a result of some unknown principles), as in the well-known Baldwin illusion or so-called line-length assimilation.

Another study of perceived spatial extent was conducted by Levi and Westheimer (1987), who determined whether variations in the luminance profile of a stimulus consisting of two bars influenced the perceived separation of the bars, and the precision with which separation could judged. Their stimulus was made up of seven unresolved lines, with an eighth added to one of the bars to alter the location of the centroid of the retinal intensity distribution for the bar. The location of the centroid of the intensity distribution did not affect thresholds for interval width discrimination but did affect the point of subjective equality (PSE) of interval width- and the perceived width of the interval was closely matched by the separation of the centroids of the retinal intensity distribution. This is similar to the data of Westheimer and McKee (1977a) for vernier acuity, in which they showed that vernier alignment for very similar stimuli depended on the centroid of the retinal intensity distribution. This suggests initially that the way in which spatial extent is assessed and relative location in a vernier task is measured have some similarities; it may be that the perceived width of an interval is determined by the location of the spatial primitives, but there is yet another source of information for interval width judgments which is not limited by contrast (unlike localising the centroids of zero-bounded response distributions, for instance).

Levi and Westheimer (1987) also determined thresholds and points of subjective equality for stimuli composed of two bars with a lower contrast bar between them. This internal bar had no effect on threshold for discrimination of the separation of the outer bars (at contrasts up to 1/7 th that of the outer outer bars). However, the perceived width of the interval was dependent on the separation of the internal bar and the outer two bars; when the internal bars and outer bars were close, the interval between the outer two was seen as narrower. If the internal bar and the outer bars were more widely separated, the interval width was overestimated. This result is directly analogous to findings of Badcock and Westheimer (1985a), who showed that a single flanking line alters the location of a target line in the same manner; ie. the lines are 'attracted' when close and 'repelled' when slightly further apart. This phenomenon (one aspect of spatial interference) is discussed at greater length in section 3.2.

The perceived relative location of a bright bar with an asymmetric orthoaxial contrast profile was determined by Toet, Smits, Nienhuis and Koenderink (1988), using interval bisection and vernier acuity tasks. For the bisection task, they found that the spatial primitive which best explained the results was dependent on the instructions given to the subject; if the subject was asked to attend to the width of the intervals between the three bars in the target, then zero-crossings in a hypothetical neural response (equivalent to a blurred second spatial derivative) were most appropriate. Without these explicit instructions being given, the relative localisation of the asymmetric bar was best explained by supposing that the visual system extracted either extrema or centroids from the hypothetical neural response to the stimulus. The significance of this result is unclear, since there have been no comparable reports of an interaction between spatial primitives and attentional, ie. top-down factors, which can be modulated by varying the instructions given.

1.6: Models of early vision

There have been many competing models of early vision. Exponents of the two broad theoretical alignments have been labelled "frequency freaks" and "feature creatures" (Krose, 1987). These two divisions reflect the differential emphasis on information in either the spatial frequency amplitude spectrum, or about the geometric (ie. spatial) properties of the stimuli. It would be fair to say that the feature creatures are (once again) in the ascendancy.

1.6.1: Frequency based models

One example of a frequency domain approach to explaining spatial discrimination is a model devised by Carlson and Klopfenstein (1985). This was not proposed as a comprehensive model of spatial discrimination, but was designed to illustrate that in principle the change in response of spatial frequency detectors to two slightly different patterns is sufficient to allow relative localisation with 'hyperacuity' precision. Their model involved a range of medium bandwidth spatial frequency detectors, each with its own non-linearity (inspired by contrast discrimination experiments) and a noise source. Detecting a change in the spatial structure of the pattern takes place by analysing the change in response of the filter most sensitive to change in the pattern- ie. the filter which exhibits the greatest differential response between two patterns. The model predicts thresholds for relative localisation tasks, considering only the changes in the spatial frequency amplitude spectrum of patterns.

1.6.2: Wilson's model

One of the most influential models of early vision has been that of Wilson and Bergen (1979). Wilson's models have always been concerned with the energy in localised spatial frequency detectors and do not merit the 'hybrid' character that he claims for them. This statement is justified below. Wilson and Bergen's (1979) models of threshold spatial vision was derived from a re-analysis of data in a previous work by Wilson (1978); Wilson (1978) measured the influence of subthreshold flanking lines on a target line's detectability at different flanking distances, repeating this for two forms of temporal modulation ('sustained' and 'transient') as well as three different eccentricities. The line-spread function (LSF) for the system with a given set of parameters describes the sensitivity of the system as a function of the separation of the target and the subthreshold flanking lines. For lines with the same contrast polarity, the flanking lines facilitate detection when they are close to the target line, and inhibit detection when slightly further away, the inhibition effect diminishing at greater distances, giving a non-monotonic function relating sensitivity to flanking distance.

Wilson used the line-spread functions from this data set to estimate the filter profile and sensitivity of two bandpass filters which best explained the line detection data. This gave him a model for spatial contrast detection, which he used to predict thresholds for detecting Difference- of- Gaussian (DOG) patterns and cosine bars. The data were not well explained by the model predictions, with evidence of both lower and higher peak frequency bandpass filters contributing to the detection of these patterns.

Consequently, Wilson and Bergen (1979) modelled the data for detection of DOG patterns of different sizes with four filters, one higher and the other lower in peak frequency to the filters in Wilson's (1978) model. Wilson and Bergen's model is based loosely upon Quick's (1974) vector magnitude model for detecting spatial patterns, in which the responses of independent detectors are combined by some magnitude function prior to detection; the magnitude function is given by

$$\text{vector magnitude sum of } (\mathbf{R}) = [\sum (R_i)^p]^{1/p} \quad \text{where } \mathbf{R} = \{ R_1 , R_2 , ... R_n \}.$$

In other words, the responses of a range of detectors are combined into a single sum, which depending on the value of p (the exponent) will range from an arithmetic mean (p=1) to peak detection (for large p). Wilson and Bergen fitted the DOG data from Wilson (1978) with four types of filter, each having different sensitivities, different space constants and different temporal response characteristics.

Responses of individual filters are computed by convolving the filter profile with the luminance profile of the stimulus, with the response being weighted by an amount which is proportional to the sensitivity of each detector. Whilst Wilson and Bergen's model has only four scales of filter, there are a larger number of filters pooled to derive the theoretical response of the system. To predict theoretical detection thresholds for the DOG patterns, they use five filters of each type, the centres of each of these being separated by 2.0 min arc.

The reason for incorporating units which are not centred under the stimulus is to improve the fit to the data- Wilson's approach is to start with the simplest possible model configuration to see if this can explain the data, and if it does not, add an extra mechanism and see how well this new configuration fits the data. The main point of Wilson's work is that a relatively small number of spatial filters generally provide adequate explanation of detection sensitivity.

Wilson's basic model was elaborated following further masking experiments. Wilson, MacFarlane and Phillips (1983) determined threshold elevation for spatially localised D6 (sixth derivative of a Gaussian) patterns, as a function of the frequency of a masking sinewave grating at 15 deg orientation to the test pattern. Wilson et al found that elevation of thresholds often did not peak at the test pattern's peak frequency, but was slightly displaced. For instance, masking of both a 1.4 and 2 c/deg pattern for one subject was optimal when the grating was of 2 c/deg. Analysis of the threshold elevation curves obtained indicated that as few as six filters were necesssary to explain the threshold evelation data, within the limits imposed by experimental error.

Wilson and Bergen's model of detection was extended to cover discrimination tasks by Wilson and Gelb (1984), in their 'modified line-element' model of spatial discrimination, named by analogy with line-element models in colour vision. Six peak frequencies of filters were employed in Wilson and Gelb's model, ranging in peak frequency from 0.8 to 16.0 c/deg. Filter sensitivities and bandwidths were estimated from the masking data of Wilson et al (1983). The response of each filter to the stimulus is computed by convolution of the pattern's luminance profile with the filter profile, the response is then weighted by the filter's sensitivity. Then the response is passed through a contrast transfer function, with each filter having a different transfer function. Response of each filter to the first pattern is then stored. The procedure is repeated when the second pattern is presented to the system.

The difference in response magnitude of each individual filter to the two patterns is then computed for each filter type; if we write the difference in magnitude of the i^{th} filter response to the two patterns as dR_i and the vector of differences as $\mathbf{dR}$

$$\mathbf{dR} = \{ dR_1, dR_2, \ldots\ldots, dR_n \} \text{ for n filters}$$

then Wilson's model involves computing the vector magnitude sum for **dR** using Quick's (1974) formulation given above. The difference in response of each filter size to each pattern is computed for filters centred under the stimulus and for one filter either side of this centred filter. The model pools across these spatially offset filters, and therefore discards potential information about the spatial distribution of responses, taking only the magnitude of the difference in response in each filter to each of the two patterns. This model therefore does not merit being called a 'hybrid' model, since there is no information preserved about the spatial distribution of filter response, unlike models such as MIRAGE (Watt & Morgan, 1985) or Marr and Hildreth (1980).

In Wilson (1986), a two-dimensional extension of Wilson and Gelb's (1984) model is applied to a wide range of pattern acuity data; this model (Wilson, 1986) is identical in form to Wilson and Gelb's model, but incorporates orientation selective filters, and pooling across all orientations before a decision is made. Wilson (1986) attempts to explain a wide range of data with a single model with very few free parameters. Wilson notes that his model is inappropriate or fails completely to explain performance on relative localisation tasks involving widely separated, relatively fine features (eg. separated by more than 1 deg). This is principally because it does not allow for activity in filters which are not spatially near-neighbours as a source of information about relative location. It cannot do so, because information about the spatial location of filters is not preserved at all in Wilson's scheme and this seems to be its basic philosophy. Very low frequency filters would have to carry information about the relative location of widely separated features, lower than Wilson could find evidence for with his masking experiments.

Wilson's model has a very large 'null-space'; this is the equivalence class of patterns which Wilson's model fails to discriminate (Nielsen & Wandell, 1988). Given the representation that Wilson advocates does not have the property of uniqueness, there are many ways of mapping different patterns into the same point in Wilson's multidimensional filter response space. Additionally, there are many patterns which would be seen as perceptually identical which Wilson's model identifies as quite different. Since Wilson's model discards information about local image geometry, it is capable of making discriminations only when the spatial frequency amplitude spectra of patterns have a meaningful relationship to the changes in pattern which have to be identified.

As empirical evidence against the approaches of Wilson and Carlson and Klopfenstein, Morgan and Ward (1985) demonstrated that accurate spatial discrimination is still possible in the presence of considerable perturbations in the spatial frequency amplitude spectra of patterns to be discriminated. Morgan and Ward note that these approaches can succeed if there is a difference in the spatial frequency amplitude spectra of two patterns, however complex these spectra may be. However, if the spatial frequency content of stimuli is randomly varied from trial to trial, then these models would predict a drastic decline in performance. Morgan and Ward performed two experiments to test this prediction; in the first, the effects of randomly scaling a four line target on interval discrimination were determined. Subjects had to determine whether the central interval of three was narrower or wider than the outer two intervals (which were of equal width); random scaling of the target dimensions did not impair accuracy on this task. This indicates that absolute spatial frequency amplitude information is not necessary for accurate discrimination, although a system that could analyse changes in the 'shape' of the amplitude spectrum (ie. the relative distribution of energy across spatial frequency) could deal with this situation. However, changes in the activity of any particular subset of filters would not give useful information about the intervals to be discriminated, which is contrary to simple models like Wilson's or Carlson and Klopfenstein's.

To eliminate the possibility that the system performs some more elaborate analysis of the frequency amplitude spectrum of each pattern, they determined interval thresholds for a target which consisted of four lines, which was compared with a two-line reference interval . The separation of the outer lines was randomly varied; the subjects had to discriminate the width of the innermost interval and the width of the two line reference interval. This manipulation will randomly perturb the 'shape' of the frequency amplitude spectra of the two patterns from trial to trial. They found that the random perturbation in flanking distance of the outer lines had not effect on threshold for the interval task. To explain this result, Wilson's model would have to conclude that the system can discover which filters carry most information about the change in interval width (rather than those which have the biggest differential response to the two patterns). Wilson (1986) mentions that a different strategy would be needed for this stimulus situation anyway, since the Wilson model would indicate that the two stimuli to be compared in this experiment are different in the first instance; but this seems to refute his own claim of having a unified theory of pattern discrimination.

1.6.3: Computational approaches : localising intensity changes

This section describes some computational approaches to the detection of intensity changes in the image. The process of identifying and assigning location to intensity discontinuities is important as a means of data reduction, and may also be critically necessary for higher-level visual processes to work dependably. Conventionally, intensity changes are detected with differential operators. Differential operators are notoriously sensitive to noise, but we can overcome this by smoothing the image; also, since we do not know in advance how blurred an intensity change in the image will be, we may have to look at the filtered image on many spatial scales; generally, signal-to-noise ratios will differ for all edges in an image, and it is necessary to incorporate several different operators with different scales in to the scheme for detecting edges (the largest filters having good sensitivity but poor localisation ability). Marr and Hildreth (1980) were the first to do this. Integrating the outputs of different scales of operators is a problem which has not yet been adequately solved.

Marr and Hildreth (1980) were also influential in proposing the use of Laplacian of a Gaussian function. Their model of edge localisation involved blurring the image with Gaussians of different space constants, and taking the Laplacian (the sum of the pure partial derivatives in orthogonal directions) of these blurred images. In practice, given the linearity of the convolution operator, we can combine these two steps into one convolution, by using an isotropic difference-of-Gaussians function, which for a centre-surround ratio of 1:1.6, gives a good engineering approximation to the Laplacian of a Gaussian. Marr and Hildreth's symbolic representation of edges consisted of a set of zero-crossing segments detected by operators with different spatial scales, along with the slope of the directional derivative perpendicular to each zero-crossing segment. Alternatively, since zeros are difficult to detect in a noisy system like the brain, they suggested that the location of zero-crossings could be estimated by interpolating between extrema of opposite sign in the filter output (Marr, Ullman & Poggio, 1979). The zero-crossing segments from the different spatial scales of filters were then combined into one representation by an algorithm based on the "spatial coincidence assumption"- if zero-crossings from filters adjacent in scale coincide, then this is enough to indicate the presence of an edge. Provided two filters that are reasonably separated in peak frequencies signal an edge, an edge primitive is represented in the 'primal sketch', a symbolic representation of intensity changes in the retinal image.

1.6.4: The optimal edge-finder

Canny (1983) proposed a computational theory of edge detection based upon Marr and Hildreth's scheme. He showed that detecting the peaks in the output of an operator which is reasonably closely approximated by a first derivative of a Gaussian is the optimal method (in a mathematical sense) of localising step edges. This is equivalent to detecting the zeros in the output of a second derivative of a Gaussian operator. He gave an algorithm for combining information from different spatial scales of operator. This algorithm entailed starting with the edges detected by the smallest operator, and using these edges to predict the outputs of larger filters. If there is an edge with a signal-to-noise ratio too low for the smallest filter to detect (since the filter sensitivity varies directly with its spatial variance) then the predicted output of larger filters will be different from the output predicted assuming that the edges detected by the smaller filter are the only edges in the image. If the difference between the 'synthetic' output and the real output of the larger filter is significant, then an edge is added on a coarser scale edge map. Canny points out that the expense in using an oriented filter is not that great, since the orientation domain can be quite roughly sampled with adequate results; eg. using 5-8 orientations of detectors seems adequate. Given that practically all cortical cells are oriented, isotropic operators may not be relevant to biological vision, although this issue is not definitely settled.

1.6.5: MIRAGE

A model of the representation of intensity changes in the early stages of the human visual system which is related to Marr and Hildreth's scheme was devised by Watt and Morgan (1984, 1985). There are some key differences in the way that primitives are extracted from the output of the filters. In Marr and Hildreth's (1980) model, each scale of filter delivers an independent estimate of the location of intensity changes in the image; if filters of different frequencies signal a zero-crossing at the same retinal locus, then an edge primitive is placed at that location (the spatial coincidence assumption). Canny's scheme uses different scales of filters to independently analyse the image, with the responses of smaller filters being used to determine whether activity in a larger scale of filter is due to an intensity change not detected by the smaller filter. With this approach, a major problem with the Marr and Hildreth algorithm is overcome, that adjacent edges of fine spatial scale will distort the output of larger filters and make the spatial coincidence assumption invalid (and shift the location of the edge primitive relative to the centre of the physical edge).

In Watt and Morgan's model, the image is sampled by independent filters as in the other schemes but the information from different scales of filter is combined before the system can analyse the output of each scale of filter. This achieves two things; the effect of uncorrelated noise in the filters is reduced, and a simple coding which preserves useful information about image structure is achieved. In MIRAGE, primitives are extracted after analysing the combined responses of a range of bandpass filter, these responses having been combined in such a way as to preserve some of the information carried by the higher-frequency filters. This is achieved simply by splitting the filter response into positive and negative parts (half-wave rectification) and summing the positive and negative parts across the different scales of filter, which results in two spatial distributions of response. Note that the outputs of the different scales are independently normalised for contrast.

This operation is followed by an analysis of these positive and negative distributions of response, resulting in an interpretation of the image and extraction of spatial primitives. Noise is used to truncate regions of the two spatial signals, and the positive and negative spatial signals are changed into a series of noise-bounded response distributions characterised in terms of the location of their centroids, standard deviation and mass. These parameters allow information about edge blur, contrast and spatial location to be easily estimated. For a step edge, the edge will lie between two centroids of opposite sign. The blur of an edge can be estimated from either the standard deviation of the region of response, or the separation of the centroids; the amplitude of the intensity change can be assessed by dividing the mass by the standard deviation. Simple rules are then applied to this characterisation of the response distribution to extract primitives of the appropriate sort. An edge primitive is marked where there are two centroids of opposite sign back-to-back, and a bar will be marked where two centroids of the same sign have a centroid of the opposite sign between them. Centroids are not grouped together if separated by an appropriately broad region of zero filter response- circa 3 min arc (see 1.5.3 above).

The major difference between MIRAGE and comparable models of early vision is the collapsing of information from different filters before the analysis of filter responses takes place. One of the critical pieces of evidence is that blur extent difference thresholds are higher for a sharp edge than for one blurred over 5 min arc (Watt & Morgan, 1983b). Watt and Morgan (1984) attributed this to addition of a moderate amount of internal blur to the representation of an edge, and concluded that the smallest filters could not be accessed independently of the larger filters, and involvement of a filter with peak frequency of about 3 c/deg was necessary to explain the data for blur discrimination at all edge blurs. Morgan and Watt (1984) demonstrated that activity in smaller filters can prevent information in the larger filters being extracted with normal efficiency. They constructed a stimulus which was sampled from a sinewave grating, which had to be aligned with a section of normal sinewave grating. The sampled stimulus has substantial energy at its fundamental frequency, and the spatial distribution of response in a large filter will reconstruct the waveform from which the stimulus has been sampled by effectively interpolating between adjacent samples. Morgan and Watt determined vernier acuity for this stimulus, and then masked the sampled grating with a high frequency grating. Presence of a mask with no frequencies lower than 40 c/deg impaired vernier acuity markedly for a target stimulus sampled from a 1.5 c/deg grating. The main point is that the output of low-frequency filters could not be accessed independently from the higher-frequency filters. This effect was phase dependent, and randomising the relative phase of the mask to the target reduced the interference with vernier acuity, indicating the masking was caused by the spatial structure of the mask rather than just energy at a high spatial frequency. A similar result was reported by Watt and Morgan (1984); they determined vernier acuity for a coarsely blurred edge in the presence of a patch of high-frequency grating at the centre of the edge, and at the sides of the edge. If the grating is placed in the centre of a blurred edge, then vernier acuity is degraded, indicating that fine detail can interfere with extraction of low frequency information.

1.7: Neurophysiology and early vision

There is evidence that cells in primary visual cortex are selectively sensitive to intensity changes in the retinal image (Hubel & Wiesel, 1962, 1968), such as bars and lines of a particular orientation. The receptive field profiles of cells in V1 bear an uncanny resemblance to filter profiles used in models of spatial vision; cortical cells can be thought of as oriented bandpass spatial frequency filters (De Valois, Albrecht & Thorell, 1982; Daugman, 1985).

1.7.1: Physiology and Function in Cortical Cells

A common technique in physiological studies of the visual system is to determine the response of a cell to sinewave gratings of different spatial frequencies; its frequency response can then be Fourier transformed to derive an estimate of the spatial weighting function of a cell. An alternative is to present spatially localised stimuli, like bars, spots and lines (Jones & Palmer, 1987; Hubel & Wiesel, 1962) and map out the spatial weighting profile of a cell directly.

Both these techniques for mapping out receptive field profiles only work if a cell is linear (or approximately so); that is, if it shows linear spatial summation across its receptive field, in which case its response can be described by convolution of the plotted receptive field profile with the luminance profile of the stimulus. Simple cells obey this condition, but other sorts of cortical cells do not; for such non-linear cells, some estimate of the cell's preferences can be obtained by determining which stimuli provoke the most vigorous firing rate, but its response cannot be described by convolution of the stimulus with any receptive field profile. Non-linear receptive fields can be mapped with other techniques, well outwith the scope of this discussion (eg. using Wiener kernels).

Spatial summation properties were originally used by Hubel and Wiesel (1962, 1959) to classify cells as 'simple' or 'complex', with simple cells showing approximately linear spatial summation in the dimension orthogonal to their preferred orientation. For a 'linear' cell, the response to superimposed stimuli is the algebraic sum of the response to the individual stimuli. Simple cells typically show a 'phase-null' position for sinewave gratings (eg. De Valois, Thorell & Albrecht, 1982) - the cell's response is reduced to its spontaneous activity level when the grating is in a certain phase relationship to the centre of its receptive field (eg. in cosine phase for an even-symmetric receptive field). For all sorts of cells, it is possible to determine which range of spatial frequencies can best drive the cell; the narrower the range of frequencies, the narrower a cell's spatial frequency bandwidth. A narrow spatial frequency bandwidth implies a cell with many sidelobes in its receptive field profile. A spatially restricted receptive field profile implies a high peak spatial frequency in the cell's response; larger receptive fields generally (but not necessarily) imply a lower peak spatial frequency to a cell's spatial frequency tuning curve.

1.7.2: Cortical cells and representation

Several functions have been proposed to describe the physiologically-determined responses of simple cells. A common candidate is the Gabor function, which has convenient mathematical properties (eg. as a basis for image coding, it gives a complete representation). Gabor, in developing a theory of communication, formally proved that a sinewave modulated with a Gaussian envelope achieves the theoretical lower limit of joint uncertainty in time and frequency (Daugman, 1985). The Gabor function has been used widely in analysing the responses of cortical cells (eg. Kulikowski, Marcelja & Bishop, 1982; Daugman, 1980). There is conflicting evidence as to how well Gabor functions described simple cell spatial sensitivity profiles. Daugman (1985) claims an excellent fit of Gabor to simple cell response profiles collected by Jones and Palmer (unpub. data, in Daugman, 1985), as do Kulikowski et al (1982). Recent neurophysiological data from Parker and Hawken (1988) suggest that the

Gabor is an adequate description of the response of only a small subset of simple cells; generally, the low spatial frequency cut-offs of Gabors (and second-derivatives of Gaussians, or D^2G's) are too steep to accurately describe the cells frequency responses. However, Difference-of-Gaussian (DOG) functions provided a fairly good fit to all of their data. It is somewhat unclear whether the response functions of cells measured in this way represent serious constraints on the sorts of filters which may be part of a mathematical model of early vision.

1.7.3: Neural models of spatial discrimination

Several studies of the responses of individual cortical neurones have demonstrated that single neurones can respond to visual stimuli with a sensitivity that is about equal to the sensitivity of the entire system. Parker and Hawken (1985) measured the response of simple cells in striate cortex of monkey to a sinewave grating; since these cells show linear spatial summation, their responses are phase-dependent. When the phase of the grating relative to the centre of the cell's receptive field is changed, the response of the cell changes. Changing the phase a small amount (eg 10 deg of phase angle) can reduce the probability of the cell giving at least one spike when the grating is presented from near 1 to very near 0. The signal-to-noise ratio of the cells was high enough to give estimated localisation thresholds of as low as 11 sec arc, well within the 'hyperacuity' range. Frequency discrimination thresholds for cells were estimated to be as low as 3.7%, in the region of human performance (eg. Campbell, Nachmias & Jukes, 1970).

One interesting observation of Parker and Hawken is that the receptive fields of the cells with the finest localisation performance were larger than estimates of the smallest channels in the visual system, based on consideration of two point and line acuity psychophysical data (Marr, Poggio & Hildreth, 1980). Rather, the receptive fields widths of the best-performing cells on a localisation task are typically in the range 4-5 min arc, which is the distance over which spatial interference with judgments like vernier acuity is greatest (Westheimer & Hauske, 1975).

One major problem with this interpretation of neurophysiological data is that a number of factors are known to affect the firing rate of cells early in the visual pathway; human psychophysical observers seem to be very good at extracting measures of pattern parameters which are invariant under changes in the pattern which would modulate the responses of large numbers of cortical cells in a way not related to the task; for instance, the results of Burbeck and Regan (1983) and Bradley and Skottun (1984), that orientation discrimination and frequency discrimination are scale- and rotation invariant respectively cannot easily be reconciled with single-cell interpretations of psychophysical performance. Further, Skottun, Bradley, Sclar, Ohzawa and Freeman (1987) measured the responses of cortical cells at different contrasts to changes in orientation and spatial frequency of grating patterns, and noted that most cells gave estimates of thresholds which were dependent on contrast for a wide part of the contrast range, whereas spatial frequency discrimination is independent of contrast over a wide part of the contrast range. However, Parker and Hawken's results indicate that signals provided by individual cells are quite reliable, and it may be possible to compute an invariant description of the change in a pattern using the responses of a relatively limited number of neurones.

These are localist models of neural operation; at the other end of the spectrum are population models of encoding, which exploit the conjoint activity of many cells. to overcome the shortcoming of using individual cells in image encoding (ie. the fact that they modulate their activity in response to changes in many parameters of the stimulus). One example of a model of this sort was developed by Paradiso (1988), to explain certain aspects of orientation discrimination, eg. spatial interference with orientation discrimination (Westheimer Shimamura & McKee, 1976). This model was based upon an individual hypercolumn, first proposed as a mechanism for the analysis of orientation by Hubel and Wiesel (1974a,b). Levi, Klein and Aitsebaomo (1985) noted that the retinal distance over which spatial interference is greatest is roughly equal to the estimated width of an orientation hypercolumn in macaque and to the width of a human ocular dominance column. They studied spatial interference with vernier acuity at different eccentricities, and concluded that it was related to the width of a cortical hypercolumn, which they suggested operated as a processing module. Paradiso's model of orientation discrimination was based on using the activity profile across many cells in an orientation hypercolumn as a representation of the orientation of a contour in the image. Orientation is assessed by comparing this activity profile with some orientation template. The model takes into account the variability and broad tuning (in several domains) of cortical cells, and explains orientation discrimination performance, some orientation illusions, and spatial interference with orientation discrimination, but using at least 1,000 cells to achieve performance equivalent to psychophysically measured thresholds.

Chapter Two

Methodology

2.1: Overview

The material in this section presents details of the psychophysical methodology used, and brief details of the computational modelling employed. The details of the apparatus used to present the stimuli, control the experiment and record responses are given in section 2.2, Sections 2.3 and 2.4 give information about the subjects used and some of the general experimental details common to all of the experiments performed. An account of the modified version of QUEST used to collect data is given in section 2.5, along with a justification of the changes made to Watson and Pelli's (1983) original version.

2.2: Apparatus

2.2.1: Direction of displacement discrimination

In experiment 3.1, a direction of displacement discrimination task, patterns were generated by computer, which was also used to control the experiment, collect data and analyse the subjects' responses. Tables for the luminance profiles of patterns were computed and passed via digital-analogue converters (DAC's) through a custom-built filter, designed to remove the spurious high spatial frequencies introduced by the digital quantisation. A raster was provided by a signal from a Tektronix 500 series function generator- onset of the patterns was synchronised with the flyback of the display by using the pulse generator on this module. Patterns were displayed on a Tektronix 606A monitor with P31 phosphor. The display was effectively linear over the luminance range employed, with a maximum deviation from Z-axis linearity of only 2%. Frame rate of the display was 110 Hz.

2.2.2: Vernier acuity and spatial interval discrimination

Patterns were generated by computer, which was also used to control experiments, collect and analyse data. The computer was connected to a Cambridge Electronics Design 1401 Intelligent Lab Interface. This 1401 Interface was in turn connected to an Innisfree PICASSO CRT Image Generator. Four DAC channels were available on the 1401 to control the image generated by the PICASSO; for instance, it was possible to spatially modulate one or two fields of luminance, control the windowing and spatial position of one field, control the orientation of the two fields. The PICASSO generated a frame synchronisation pulse which triggered the output of the 1401 at the start of each frame to be displayed.

The CRT outputs from the PICASSO drove a Tektronix 606A monitor, with P31 phosphor, the same monitor as described above, with the difference that it was driven at a 100 Hz frame rate. Data was collected with a Tektronix 608 monitor in a few of the experiments, with all other equipment being the same.

For all experiments, patterns were digitised to 512 pixels in the direction of luminance modulation. For all but the first series of experiments, the patterns were displayed in a field of 77 mm. Mean luminance of the display was 22.6 cd/m^2 for the 606A, with a maximum stimulus contrast of 0.45 (Michelson contrast).

2.3: Subjects

Subjects from whom data was collected were, whenever possible, psychophysically experienced. With the exception of the author (IRP), all subjects were naive to the purposes of the experiments. Extensive practice was given before data collection began (ie. some 2,000-3,000 trials or more, typically). All of the observers, with the exception of one (whose refractive error was less than 1.0 D and wore spectacles for the data collection) wore contact lenses, which minimise changes in magnification of the retinal image compared to spectacles, and which avoid the problems of spectacles such as changes in the field of view and off-axis aberrations. None of the observers had errors greater than -3.5 D uncorrected and none had astigmatism greater than 0.5 D.

2.4: General experimental details

2.4.1: Experimental sessions

Subjects were invariably allowed practice trials at the start of an experimental session. This practice session lasted for some 3-5 minutes on average, and served the purposes of adapting the subjects to the screen luminance, as well as re-orienting them with the experiment. An experimental session consisted of some 600-900 trials, broken up into blocks of 120-200 trials, with each block lasting some 5-10 minutes on average. An instant before each stimulus presentation, a tone was provided by the computer to alert the subject to the stimulus onset. For those experiments which used a temporal two-alternative forced-choice procedure, a standard inter-stimulus interval (ISI) of 500 msec was employed.

2.4.2: Control of fixation

For experiment 3.1 only, no fixation point was employed. Subjects were instructed to maintain fixation upon the centre of the display as best they could. This experiment involved judgments about the direction of displacement of a target, and it was felt that a fixation spot would have provided an additional cue to the direction of displacement. In all other experiments, fixation was controlled by asking the subjects to fixate upon a very small point in the centre of the display. This was made as small as possible whilst still remaining visible. Pilot trials were run to check that this did not artefactually reduce the spatial thresholds obtained (eg. in vernier tasks) by providing an alternative spatial reference for target localisation.

2.4.3: Spatial jitter of patterns

For all of the spatial interval experiments a random spatial displacement was added to the location of each pattern. This was to prevent a possible spatial displacement cue being present to aid judgments about the pattern, as well as to minimise a possible cue from after-images.

2.4.4: Collection of responses

At the end of each trial, the subject transmitted their decision to the computer by means of an infra-red response link or a electro-mechanical response box, these being connected to the digital input of the CED 1401 Lab Interface.

Subjects were restricted to making binary decisions about the patterns presented, or requesting a repeat presentation. Repeat presentations were encouraged only when the subject had failed to fixate immediately prior to presentation of one or both of the patterns presented. In the case of temporal 2AFC, subjects were asked to indicate in which interval, for instance, the wider stimulus had occurred. In vernier tasks, which were presented as a one-interval 2 AFC, the subjects made a decision about whether the lower half of the target was offset to the left or right of the upper half. Each of the experiments was self-paced, with the subject able to take a short break by witholding their response for a brief period. At the end of each block of 120-200 trials, the subject was allowed a short break to reduce fatigue.

2.5: Psychophysical procedure

A modified version of QUEST (Watson & Pelli, 1983) was used to select the patterns to be displayed on each trial, and is described below. Probit analysis (Finney, 1971) was used to determine the mean and standard deviation of the subject's response distribution at the end of the experiment. In practice, the procedure was configured to behave very similarly to Watt and Andrew's (1981) APE.

2.5.1: QUEST

A version of QUEST (Watson & Pelli, 1983) was used to collect data and select the stimulus level to be presented in the course of an experimental session. This version of QUEST was modified to incorporate a number of recommendations made by several authors, as well as features which made it closely resemble Watt and Andrew's (1981) APE procedure.

QUEST is a conceptually elegant procedure, based on Bayesian statistics, which makes efficient use of prior knowledge about the probable location of threshold, as well as all of the data collected in the run so far. The prior knowledge used to determine the initial testing level is not used in calculating the final estimate of threshold, this depends only on the data set (Watson & Pelli,1983).

QUEST works by computing a log-likelihood function which is updated after each trial. The modal value of this function is the current best estimate of threshold. Convergence of the procedure is achieved because after a number of trials, additional trials have less and less effect on the position of this modal value in the function, and it is possible to terminate the procedure when the variance of the log-likelihood function is sufficiently low (or its inverse variance sufficiently high), which was an amendment suggested by Laming and Marsh (1986).

Rather than using the 75% point in the psychometric function as threshold, Watson and Pelli (1983) recommended a different point; this Laming and Marsh (1986) called the sweet point of the psychometric function, which is the intensity value which makes the largest contribution to the inverse variance of the QUEST log-likelihood function. This corresponds to the 94% correct point if a cumulative normal is used to model the psychometric function. In practice as discussed below, a 91% convergence point was employed.

The present study modelled the subject's response distribution with a cumulative normal distribution. Laming and Marsh (1986) demonstrated that QUEST yields bias-free estimates of threshold with this form of the psychometric function. The amendments incorporated into QUEST are described below.

2.5.2: Termination criterion

It was not thought wise to adopt the confidence-level termination rule offered by Watson and Pelli (1983), or the termination rule based on the inverse variance of the log- likelihood QUEST function recommended by Laming and Marsh (1986). A recent study by Madigan and Williams (1987) reported that the termination criterion is unreliable and can result in premature termination of the procedure if the initial estimate of the slope of the psychometric function is too shallow. A similar observation was made by Laming and Marsh in regard of their own termination criterion. Madigan and Williams suggested running QUEST for a fixed number of trials to overcome this potential difficulty. This suggestion was adopted; QUEST was run with between 120 and 160 trials per complete psychometric function, but always for a fixed number of trials, this decided in advance of the experimental session. In practice, it was found necessary to carry out more trials when the initial estimate of threshold was much too low, but generally the procedure was robust to inaccurate initial estimates of the probable location of threshold.

2.5.3: Double QUEST

Since many of the experiments performed required an estimate of the mean of the psychometric function, it was necessary to adapt the manner in which QUEST collected data to enable this. For example, to determine a vernier threshold for the relative location of two bars, the following procedure was employed. Two randomly interleaved QUEST procedures would be conducted, one determining the vernier threshold for offsets of the bottom bar to the left, and another for vernier offsets to the right. At the final stage of data analysis, the data from these two condition collected by each individual QUEST procedure would be combined. This was done, for the vernier task, by the following rule; data was converted from percentage of response correct, to percentage of response which had been "bottom bar to right".

This method of combining data from two QUEST procedures run simultaneously and randomly interleaved transforms the response distribution from one running from 0.5 to 1.0, to one running from 0.0 to 1.0. It hence resembles the response distribution collected by the APE procedure, which also allows estimation of the mean in addition to threshold, unlike PEST (eg. Pentland, 1980) or the original QUEST.

2.5.4: Dispersed sampling

As Watt and Andrews (1981) observe, it is not desirable to present a long sequence of stimulus level of closely similar intensity, which procedures like PEST and QUEST very often will do. This can produce a serial correlation in the subject's responses, ie. a serial response bias, which vitiates the measurements taken. One advantage of APE over these other procedures is that stimulus levels are presented in a pseudo- random sequence. Watson and Pelli (1983) note that this may overcome subject boredom and frustration at being presented with many trials at a threshold intensity.

Additionally, since it was necessary to obtain a good estimate of the slope of the psychometric function, QUEST was configured to select stimulus levels of widely differing difficulties, as in APE. For reasons discussed later, QUEST was set to converge to the "sweet point" of the psychometric function, equivalent to about 94% correct. Since in practice subjects are liable to make "finger blinks" (Watson & Pelli, 1983), a convergence level of 91% was chosen, to allow some leeway for errors due to inattention. The procedure also sampled at a stimulus intensity half the intensity of convergence level, which corresponds to about 76% correct.

As in APE, the mean of the psychometric function, as well as an initial estimate of the sweet point could be supplied by the experimenter. This was necessary to ensure well-behaved sampling for some of the experiments where a shift in the mean of two or more times the spatial threshold was often obtained.

2.5.5: Probit analysis

At the recommendation of Watson and Pelli (1983), data collected by QUEST was subjected to a conventional maximum likelihood estimation of the parameters of the response distribution. This was done using the method of probits (Finney, 1971). The method of probits allowed an estimate of the standard deviation (taken as 1 JND, = 84% correct) and the mean (50% point) of the of the response distribution. The standard deviation of the response distribution is an estimate of subject performance on a task; the mean of the response distribution is the point of subjective alignment (for a vernier task) or point of subjective equality (in an interval width discrimination task).

2.5.6: Multiple interleaved QUESTs

In many of the experiments, it was necessary to randomly interleave many different stimulus conditions to reduce possible response biases by the subjects (ie. expts. 1-4). Up to 16 separate QUEST functions were simultaneously maintained by the computer, allowing up to 8 threshold and mean determinations from one single experimental session. Additionally, the computer program saved the state of the experiment after each block of trials, so that subjects could perform small blocks of trials in the highly interleaved condition and avoid fatigue. Typically, between 120 and 200 trials would be performed in blocks lasting from 5-8 minutes on average. A complete experimental session would involve between 600 and 960 trials.

2.5.7: Convergence point

Choice of the sweet point recommended by Laming and Marsh (1986) as the convergence level for QUEST was borne out by the reliability of the data obtained. The procedure could also be configured to converge to other points on the psychometric function; eg. the 84% correct point was used in some studies, which allowed greater tolerance to variable performance. The convergence of the procedure was checked by running a large number of trials, ie. over 250 for one QUEST procedure, and establishing the percentage correct at the terminal QUEST value. This was invariably very close to the level of convergence chosen at the start of the procedure.

2.6: Modelling of localisation of luminance changes

The primary motivation of the simple modelling performed was to attempt to explain the results obtained in experiments involving judgments about the relative location of luminance changes, using simple one-dimensional models with as few free parameters as possible. The models used are based on previous empirical and theoretical work (Toet, Smits, Nienhuis & Koenderink, 1988; Watt &

Morgan, 1985). If these models can account for previously unexplained non-monotonicities in perception in a straightforward way, then this is of theoretical interest.

2.6.1: Stimuli

Modelling was performed on a VAX 11/785 computer; some additional models were run on a MacIntosh 11cx, with the Mathematica software package. The orthoaxial contrast profile of the stimulus was quantised to a maximum resolution of 1.9 seconds of arc, representing 1024 luminance samples in a field with minimum width of 32 min arc (for the finest stimuli employed) with the field width being scaled for larger stimuli.

The contrast profile of the stimulus was then convolved with an approximation to the optical line spread function; for a 2mm pupil, this is close to a Gaussian with a space constant of 22 sec arc (Cambell & Gubisch, 1966). Convolution was performed by multiplying together the Fourier transforms of the functions to be convolved and inverse- transforming. An early non-linearity in transformation of retinal intensity into retinal ganglion cell responses was incorporated (h-transform; see chapter 4), using identical parameters to Toet, Smits, Nienhuis and Koenderink, 1988. The optically blurred (and h-transformed) stimulus was then transformed into a hypothetical neural response distribution; in the simplest case, this corresponded to convolution with a DOG function (Toet, Smits, Nienhuis & Koenderink, 1988; Marr & Hildreth, 1980). MIRAGE (Watt & Morgan, 1985) was implemented largely as recommended by Watt and Morgan. For MIRAGE, the filters employed had a DOG filter profile, with a centre-surround ratio of 1:1.6, which is a close approximation to a Laplacian of a Gaussian. Details of MIRAGE are presented in section 1.6 of Chapter one. Filter space constants ranged from 0.35-5.6 min arc.

For both the spatial interval and the spatial interference experiments, one-dimensional modelling was performed. This consisted in determining either the predicted perceived separation of two luminance features (for the interval discrimination experiments) or the changes in perceived location of one feature as a consequence of the presence of another nearby feature (for the spatial interference experiments) Previous empirical work (Watt & Morgan, 1983a,b) has demonstrated that spatial primitives can explain aspects of perception which other sorts of models simply do not address. Drawing upon this work, it is possible to devise many simple accounts of how relative perceptual location might be determined in the visual system. In addition, some predictions were made about the contribution of localisation of spatial primitives to response variance in situations where the stimulus configuration was subject to random spatial perturbations (eg. experiments 2 and 3, chapter 5).

Chapter Three

Spatial interference with relative localisation

3.1: Overview

This chapter is concerned with spatial interactions between closely spaced features or 'spatial interference', and whether these interactions can be explained in terms of theories about spatial primitives. Spatial interference or 'crowding' is a frequently observed phenomenon in visual experiments, and has been reported for tasks such as stereoacuity, orientation discrimination, vernier acuity, resolution acuity, and letter recognition. Since spatial interference can be obtained with dichoptic presentation (eg. Westheimer & Hauske, 1975), its origin is not due solely to retinal interactions. Some aspects of spatial interference- for instance, the manner in which it scales with eccentricity (Levi, Klein & Aitsebaomo, 1985) - suggest that it reflects a basic principle of cortical processing of relative location of image features, but at present there is no clear understanding of why it occurs. One suggestion is that it is a side-effect of localisation of intensity changes (Watt, 1984). The experiments and modelling in this section test this idea.

It was found possible to reliably measure shifts in perceived location without decreases in the precision of relative localisation. Several new observations were made. The attraction component of spatial interference is found to disappear between resolved features if the features are blurred beyond a criterial blur, and there are differences in the extent of the zone of attraction for vernier acuity and direction of displacement discrimination task. Edge localisation is affected differently to bar localisation by the polarity of contrast of flanking feature- at small separations, a bright edge target is attracted to a dark bar or edge, whilst a bright bar is attracted to a bright flanking feature. Perceived width of spatial intervals is influenced by the contrast polarity of flanking features in a manner similar to interference with vernier acuity. MIRAGE explains several aspects of the data in elegant fashion. The results imply information is combined across spatial scales, and the results are explained by extracting and representing MIRAGE-like spatial primitives. Identical parameters to those used by Watt and Morgan (1984) fit the data presented here, ie. filters with space constants from 0.35-2.8 min arc. Modelling of the data suggests an octave shift downwards in the frequency range of filters being employed in a displacement discrimination task. Data collected at varying exposure durations suggested that a coarse-to-fine strategy recommended by Watt (1987) is not used by the visual system for a vernier localisation task, at least for exposure durations of greater than 60 msec. The observation that the perceived location of a positive contrast edge is biased towards a negative contrast flank feature is explained by MIRAGE, by supposing that edge location is determined by interpolation between adjacent corresponding centroids of opposite sense.

3.2: Introduction

Shifts in perceived location of part of a figure following inspection of another "inducing" figure have been reported many times, first by the Gestalt psychologists Kohler and Wallach under the rubric of "figural after-effects" (FAE's). Kohler and Wallach (1944) observed many instances of changes in perceived size and depth of one geometric figure (the test figure) caused by viewing another geometric figure (inspection or inducing figure) for a period of time, which could be as brief as 0.5 sec. They found also a displacement effect, where the location of a test figure is displaced perceptually from its veridical relative location by inspecting an inducing figure. Typical stimulus arrangements used by Kohler and Wallach to demonstrate FAE's are seen below.

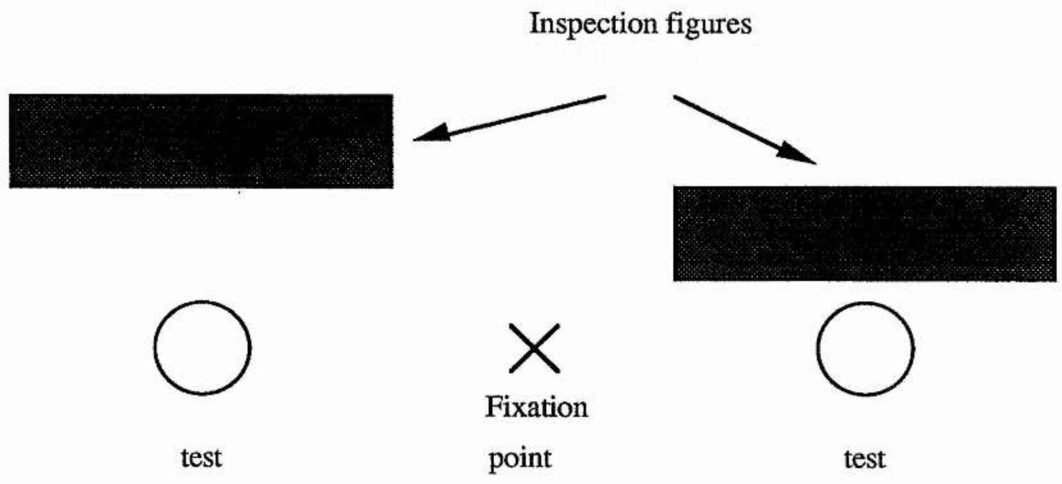

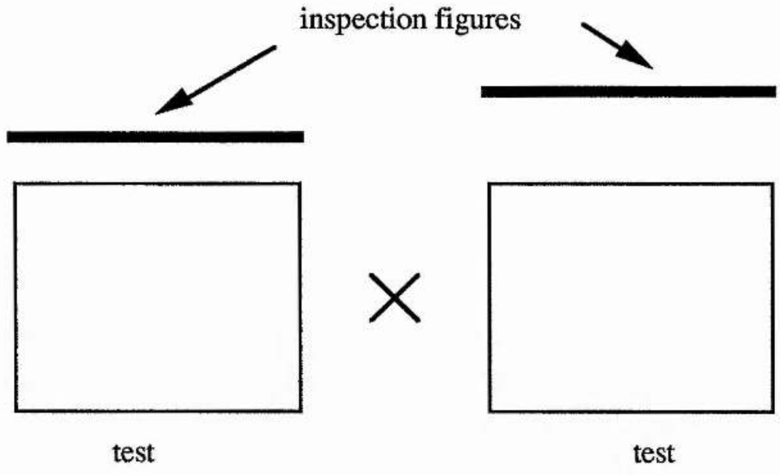

Figure 3.1: stimuli used by Kohler and Wallach (1944) to obtain distortions in perceived relative location.

When the inspection figure and the test figure were nearly adjacent, the inspection figure had less influence on the perceived location of the test figure than when the two were a little more widely separated. The inducing figure had a maximal influence on the location of the test figure when some 6-8 min arc distant from it, rather than when nearer to it, which they termed the "distance paradox". They also knew that the retinal distance at which the effect of the inducing figure on the test figure's location was greatest increased with eccentricity. They observed only shifts of the test figure away from the inspection figure, a sort of perceptual 'repulsion' between the two figures.

Later researchers employing more precise control of fixation found that perceptual attraction could occur at small inter-featural distances; the first reported attraction was by Ganz and Day, (1962) ; they found attraction between test and inducing figure at retinal distances of less than about 10 min arc (with features of the same contrast polarity), with repulsion at greater distances. From physiological studies of Ratliff and co-workers, it was already known that neural mechanisms with centre-surround antagonistic receptive field organisation existed in biological visual systems, and Ganz and Day suggested that figural after-effects were a means of contour enhancement resulting from the application of such neural operators to the visual scene. An extensive study of the FAE's was subsequently conducted by Ganz (1964), who tested also the 'simultaneous' illusion, presenting both the test and inspection figure at the same time. The stimuli used by Ganz can be seen in figure 3.2.

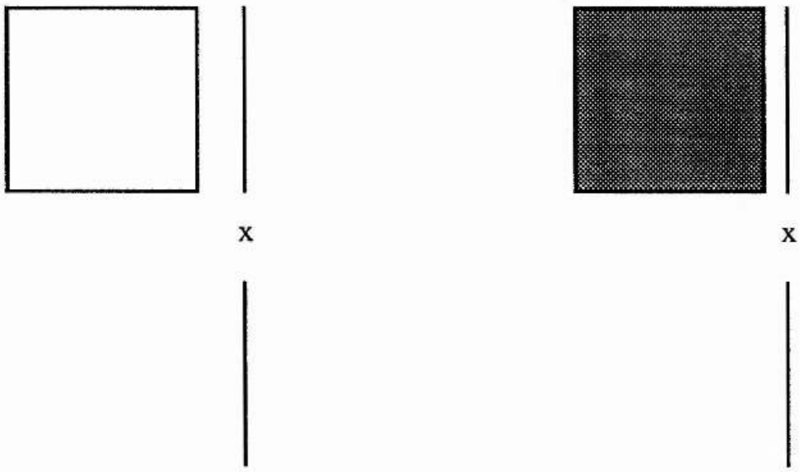

Figure 3.2: stimuli of Ganz (1964), used to quantify the influence of a flanking figure on one line in a vernier target.

Ganz (1964) measured shifts in the perceived location of a test line in normal and stabilised viewing conditions, with both successive and simultaneous presentation of inspection figure (or 'flank' for the simultaneous presentation) and test figure (or target). When a positive contrast line is placed next to a positive contrast flanking square, with a retinal distance of less than about 5 min arc between the features, perceptual attraction occurs between the flank and the target line, the line's perceived location being biased towards that of the square. Ganz considered that it was the edge of the square which was significant (rather than the entire shape). At greater retinal distance between the target and flank, only repulsion effects can be measured. When a positive contrast line is placed next to a negative contrast edge, then a repulsion effect is generated at all separations of target and flank. Ganz suggested that the edge (or its afterimage) induced lateral inhibition and facilitation in its vicinity, and generated Mach bands which interfered with perception of the target line. On the other hand, Mach bands are not seen at a sharp edge (Ratliff, 1984), but spatial interactions are observed between sharp edges and lines (Ganz, 1964). Also, Ganz's model predicts only repulsion, and no attraction.

Besides altering the perceived location of image features, nearby interfering features can also degrade spatial localisation accuracy. Recent work on the localisation of intensity changes in human vision (Watt & Morgan, 1984, 1983a,b; Toet, Smits, Nienhuis & Koenderink, 1988) has used relative localisation tasks to determine the nature of spatial primitives and the way in which these are used in assigning relative location to intensity changes in the retinal image. It seems likely on a priori grounds that a process which interferes with the normal localisation of intensity changes may be able to enhance understanding of how relative location of image features is determined, which is still not thoroughly understood. Many aspects of spatial localisation can explained by supposing the visual system assigns perceptual location to zero-crossings, centroids or extrema in a neural response distribution which is akin to a blurred second spatial derivative of the image intensity profile. Watt and Morgan (1983a) showed that discrimination of the width of a narrow bright bar with an asymmetric luminance profile can be explained by supposing that the visual system computes the separation of zero-crossings in the output of a second differential filter. They showed further (Watt & Morgan, 1983b) that data from edge location and edge blur discrimination tasks could be correctly predicted by models supposing that the visual system localises centroids of zero-bounded regions in the response of blurred second differential filters. The processes required to extract these spatial primitives will be disrupted by the presence of nearby contours and it may be that an explanation for spatial interference can be derived from these theories (eg. Watt, 1984).

Spatial interference involves either an increased difficulty of localising a target or a shift in relative location of parts of the target induced by the presence of irrelevant flanking features. In a vernier task for instance, the vernier threshold is increased in the presence of flanking features. Westheimer and Hauske (1975) measured vernier thresholds for fine lines, with irrelevant flanking features; the pair of flanking lines were disposed either parallel or perpendicular to the main axis of the target lines. The stimuli they used are shown below in schematic form.

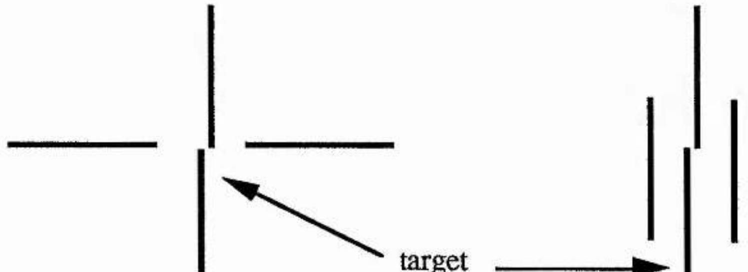

Figure 3.3: stimuli used by Westheimer and Hauske (1975), in measuring the decrease in precision of vernier alignment produced by the presence of flanking features.

The results for both configurations were quite similar, and indicated that spatial interference did not require like-oriented contours. Vernier thresholds were maximally elevated at some 3-5 min arc flanking distance. A similar spatial interference is seen in other relative localisation tasks; such as orientation discrimination (Westheimer, Shimamura & McKee, 1976) over comparable retinal distances. The neural level at which these interactions arise in the visual system is not known for certain. Westheimer and Hauske confirmed that spatial interference with vernier acuity occurs at a more central level than the retina, by dichoptic presentation of target and flanking features (ie. flanking lines to one eye and target lines to the other); this produces an elevation of vernier thresholds comparable to normal binocular presentation, indicating that the neural interactions subserving spatial interference occur at or beyond the level of binocular convergence, ie. in cortex (eg. complex cells in V1 are known to be binocular; Poggio, Motter, Squatrito & Trotter, 1985).

Further studies of spatial interference with vernier acuity shed light on the possible neural substrate of the effect. Levi, Klein and Aitsebaomo (1985) measured interference with two-dot vernier acuity for flanked dot targets, with the flanking dots in a variety of position; illustrated below are stimuli used in this experiment.

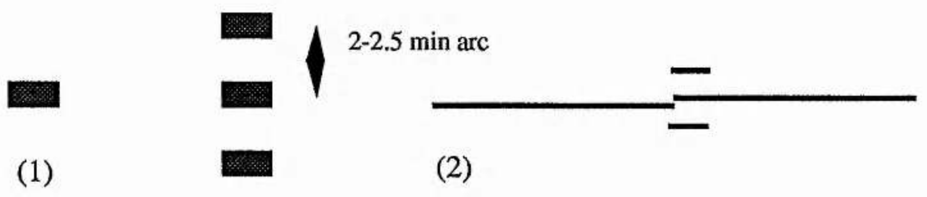

Figure 3.4: stimuli of Levi, Klein and Aitsebaomo (1985);
(1) dot stimuli used in mapping out magnitude of spatial interference ,
(2) stimuli used to measure extent of interference as function of eccentricity.

One purpose of their experiment was to map out how the relative influence flank dots had on localising the target dots varied across a small region of visual space. What they found was that the elevation in two dot vernier thresholds were greatest when the dots were aligned with the reference stimulus and about 2 min arc either side of it. Levi et al also investigated spatial interference as a function of eccentricity. As the eccentricity at which the stimulus is presented is increased, the retinal distance over which spatial interference (ie. elevation of vernier threshold) is greatest increases. For instance, at an eccentricity of 10 deg, spatial interference is greatest at retinal distances of about 30 min arc, compared with about 3-4 min arc for foveally presented stimuli.

They found that spatial interference at all eccentricities was greatest at retinal distances some 25-30 times the unflanked vernier threshold. They suggest that their data indicate the presence of processing modules, which may reflect the retinal distance over which orthoaxial contrast information is utilised in deriving estimates of the location of intensity changes. To establish a putative neural substrate for spatial interference, they noted that the unflanked vernier threshold scales closely with estimates of M, the cortical magnification factor for human striate cortex (degrees of visual angle represented per mm of striate cortex). It follows from this that the cortical distance between the neural representations of flanking and target features at which spatial interference is greatest is constant at all eccentricities, which Levi et al note is equivalent to roughly the esimated cortical width of human ocular dominance columns.

It is also known that spatial interference is influenced by the temporal modulation of the stimulus; Watt and Morgan (Watt, pers. comm.) measured interference with orientation discrimination; at exposure durations of 10 msec, the retinal distance which generates maximum spatial interference is scaled by a factor of four relative to longer exposure durations. The stimuli presented by Watt and Morgan were followed immediately by a noise mask. Watt (1987) used such masks to uncover previously unreported interactions between the spatial scale of filters involved in relative localisation tasks and exposure duration.

There are probably at least two components to the neural interactions underlying spatial interference; it is known that the attraction and repulsion zones have different temporal response characteristics, as well as different responses to contrast polarity. (Badcock & Westheimer, 1985a). Badcock and Westheimer determined vernier acuity and direction of displacement discrimination for flanked lines. For both sorts of tasks, small separations produced attraction if target and flank were of the same contrast polarity, and repulsion if of opposite polarity. Large separations of target and flank always produce a repulsion effect, regardless of the respective polarities of these features. Also, Badcock and Westheimer varied exposure duration of the stimuli, and measured the magnitude of the mean-shifts (change in perceived relative location) induced by a flank bar either 1.2 or 6 min arc from the target bar. With a flank in the 'centre' zone (1.2 min arc), the size of the mean shift increased roughly linearly with exposure duration. With a flank placed in the 'surround' zone (6 min arc) the mean-shift increased rapidly with exposure duration, and levelled at a much shorter duration than with a flank in the centre zone.

The surround rectification demonstrated by Badcock and Westheimer (ie. bright and dark flanks both inducing repulsion at flanking distances beyond 4 min arc or so) indicates that the repulsion effect is not simply a by-product of neural processing by receptive fields with luminance weighting functions with an excitatory centre and an inhibitory surround (eg. Hines, 1976; Wilson, 1978). The mechanisms subserving this effect are necessarily non-linear: that is, we cannot describe their response by a convolution. Since perceptual shifts can be obtained without degrading the accuracy of vernier acuity, it is evident that the neural mechanisms mediating the discrimination of relative location in vernier tasks can perform equally well despite distortions in the representation of image features.

In a second experiment, Badcock and Westheimer (1985b) showed that the two-dimensional spatial configuration of a stimulus affects the nature of spatial interference. A high degree of overlap in the orthoaxial direction is necessary between the target and flanking features for attraction to occur at small flanking distances. With a vertical separation between two halves of a flanking bar (ie. the flank bar not overlapping orthoaxially with the target bar), repulsion is seen at small flanking distances.

Wilson (1986) suggested that distortions in the outputs of bandpass spatial frequency filters like those in his model are the basis of spatial interference. Watt (1984) explained the degradation of vernier acuity seen in Westheimer and Hauske's (1975) experiment as due to a shift in the locations of zero- crossings in the output of a second differential filter with a space constant of about 2.0 min arc. Marr and Hildreth (1980) noted that if two features are very close together, then the smallest filters will overestimate the distance between the features. MIRAGE has already been used to explain the appearance of closely spaced luminance features; Watt and Morgan (1985) noted that the primitives generated by their MIRAGE operation could explain data on the phenomenology of the Chevreuil illusion (where an illusory third edge is seen between two step edges of the same 'sense').

Badcock and Westheimer (1985a) employed both direction of displacement discrimination and vernier acuity, and found relatively similar patterns of spatial interference with the two tasks. Recent estimates of the receptive fields sizes of human motion detection units give fields widths down to 2 min arc (Anderson & Burr, 1987). This is in line with estimates of the receptive fields widths of smallest neural mechanisms in human vision; one estimate puts the finest receptive field centres at 1'20" across (that is, the separation of the zero-crossings if the filter is represented as a Laplacian of a Gaussian: Marr, Poggio & Hildreth, 1980), from a consideration of two point and line acuity. The present experiments examine the effect that spatial scale and target/ flank identity has on spatial interference with displacement discrimination and vernier acuity. There is evidence that vernier acuity and discrimination of direction of displacement are processed somewhat differently. For instance, Nakayama and Silverman (1985) reported that direction of displacement discrimination for sine- wave gratings is a function of contrast, but that this function saturates at contrasts as low as 2-3%. Vernier acuity for sine wave gratings (Bradley & Skottun, 1987) is a function of contrast for contrasts to as high as 80%.

For blurred edges, the results are somewhat different. Mather (1987) measured displacement thresholds for cumulative Gaussian edges of different blurs, at several different contrasts. Thresholds for step displacement decrease as contrast is increased up to contrasts at least as high as 30%. Thresholds were approximately linearly dependent on the contrast over the range of contrasts tested. On the other hand, it is known that vernier acuity for blurred bars (Krauskopf & Campbell, unpublished data; Watt & Morgan, 1983b) and blurred edges (Watt & Morgan, 1984) follows a square-root relation to contrast. On a priori grounds, a difference between the effects of flanking features on displacement discrimination and vernier acuity would be expected.

3.2.1: Aims and motivation for the experiments conducted

The present set of experiments seeks to shed light on the neural substrates of localisation of intensity discontinuities, by examining spatial interference with tasks involving both bars and edges of varying degrees of blur.

The first experiment involved measuring biases in the location of a Gaussian bar by determining the 50% point which 'nulls' its apparent motion upon being given a step displacement (Morgan, Mather, Moulden & Watt, 1984, used a similar technique). The difference in the 50% point for a flanked and an unflanked bar represents the change in perceived location induced by the presence of the flanking bar. Several different blurs of bar were tested, with many different flanking distances being interleaved in any experimental session. The second experiment employs a vernier acuity task, with Gaussian bars as targets, for several blurs of bar. Changes in the point of subjective alignment induced by flanking bars were measured. Several experiments (3-5) tested the notion that localisation of bars and edges would be differentially affected by flanking features. Bar targets- edge flank, edge targets- bar flank and edge-edge combinations were tested, along with positive and negative contrast flanking features.

Recent studies of interval width perception have produced apparently conflicting results; Morgan and Ward (1985) examined interval width perception in the presence of flanking lines randomly positioned parallel to the target lines, and found no effect of the flank lines on threshold, even though flank and target lines were of the same contrast. Vernier acuity would likely have been impaired under such conditions, given the result of Westheimer and Hauske (1975) described above. Levi and Westheimer (1987) demonstrated shifts in the perceived width of a spatial interval caused by a line between two target lines, directly analogous to the results of Badcock and Westheimer (1985a). Since the flanks in Morgan and Ward's experiment were outside the spatial interval, this may allow the task to be performed with less interference than if an irrelevant line is placed in the middle of the interval between the target lines. The sixth experiment in this chapter involved spatial interference with interval width perception. Some of the results are modelled numerically, testing the predictions of Watt and Morgan's MIRAGE and single-filter models, which suppose that the visual system has independent access to the outputs of filters of different spatial scales.

3.3: Method

3.3.1: Design

Vernier and spatial interval thresholds and points of subjective alignment (equality) were determined for stimuli composed of Gaussian bars and edges of differing spatial scales. The order in which the data was collected was not systematic.

3.3.2: Subjects

The majority of data was collected from GLC, MEW and IRP, the author. Some data was collected from SM, LD and FM. All with exception of the author were undergraduates paid for participation in the study, and were naive to the purposes of the experiment. All subjects wore their normal correction (if necessary); the main subjects were highly experienced in vernier acuity and had substantial practice prior to collection of the data reported here.

3.3.3 :Apparatus

As described in Methodology chapter.

3.3.4: Control of artefacts

Since this set of experiments is measuring biases in perception, the methodology is highly sensitive to observer bias, and this is particularly problematic when the subject is not naive to the purposes of the experiment, ie. the author, IRP. For this reason, IRP was always run before any other subjects, to determine whether the relative location biases were genuine and not simply conditioned by expectation. Attempts to minimise systematically biased responding on the part of all subjects included;

(1) having an unflanked stimulus in the sequence of stimuli, to 'remind' the subject of the appearance of a genuine vernier offset.

(2) Interleaving stimuli with many different flanking distances was performed (Badcock & Westheimer, 1985a); some distances produce perceptual attraction and others perceptual repulsion, and with naive subjects, this procedure will minimise the possibility that they respond systematically (for instance) that the offset is in the direction of the flank.

(3) Having the side of the flank randomised (by a pseudo-random number generated at run-time) reduced the likelihood of systematic bias by the subject; they are less likely to develop a 'button preference' with this method. For the displacement discrimination task, the initial side of the flank was randomly

determined, with it switching to the opposite side at the instant of the jump. For the vernier bar targets with two bar flanks, the flanks in the two halves of the display were on opposite sides, but with the exact disposition of the flanks being randomised. More detail of this is available below. For the composite bar/ edge and edge-edge stimuli, the single flank was either in the top half or the lower half of the display but always on the same side of the target.

3.3.5: Expt. 1: Direction of displacement discrimination: Stimuli

Patterns consisted of vertically oriented Gaussian bars of equal width, with one (the flank) 25% of the contrast of the other (target) bar. This figure was chosen on the basis of previous studies (eg. Badcock & Westheimer, 1985a,b: Levi & Westheimer, 1987). These studies reported shifts in the point of subjective alignment without increases in the spatial threshold, and used flanking bars of lower than 25% the contrast of the target bars. Since the task of the subject is essentially to report the direction of apparent motion of the target bar, if the target and flank were not readily distinguishable on the basis of their contrasts, subjects would fail to see the target bar as moving through a small displacement, but rather the bar pair would be seen to laterally jump through a large distance. Solving the correspondence problem in motion (which object goes where across time) is an important part of the visual system's analysis of time varying images (Marr, 1982). For the stimuli with the reduced contrast flank, the subject sees the target bar shift smoothly in one direction or the other, and the flank bar is generally seen to make a large jump 'through' the target bar. Subjects were told to ignore anything to do with the flanking bar as best they could.

Another reason for keeping the contrast of the flank below that of the target is that with flank and targets of equal contrasts, an increase in vernier threshold is observed (Westheimer & Hauske, 1975), and a similar result might be expected for motion thresholds. Since the experiment initially sought to measure biases without changes in the precision of judgments, this was considered undesirable. Unlike the stimuli of Badcock and Westheimer (1985), which had the flanking line introduced at the instant of displacement, the target bar was continuously flanked in the present experiment.

Note that the data collected here measures the influence of two flanking bars on the relative location of the target bars. The flank bar changes side at the instant of the step displacement in the target bar's position. If, for example, the effect of the flank is to attract the target bar away from its veridical location before the displacement, and the target bar is moved in the opposite direction to this, whilst the bar is switched in side, then the bar will be subsequently attracted by the newly positioned flank, and the displacement will be visible at a lower physical displacement. The reverse applies if the target bar had been repelled by the flank bar in the above example. Having the flanking bar present in both intervals (but changing side) has the effect of amplifying the perceptual distortion introduced by the flank. This was verified informally in pilot trials, with greater shifts in the point of subjective alignment being induced by two flanks.

Luminance profiles for the stimuli are given in the Appendix at the end of this chapter. For any stimulus, the target and flanking bars were separated by a minimum of three times the space constant of each bar and could easily be resolved. Horizontal displacements were generated by directly shifting the raster of the display; at the viewing distance employed, a 10 bit control over position was available within a 30 min arc field, giving a theoretical spatial resolution of 0.3 sec arc.

3.3.6: Expt. 2: Vernier acuity for Gaussian bars flanked by bars: Stimuli

Patterns were composed of positive contrast polarity bars with Gausian spatial contrast profiles. Each pattern was made up of four such bars of equal width, two designated 'target' bars and two 'flanking' bars. The upper target bar was flanked on the opposite side to the lower target bar. The net effect of this arrangement would be to double the size of the observed mean shift caused by the spatial interaction between the target and flanking bars. Target bars were four times the contrast of the flanking bars. The target bars were presented at the maximum contrast available (see Methodology). The lower bar pair was given an offset relative to the top bar pair to either the left or right on each trial. The subject had to indicate the direction of this offset. Luminance profiles for the stimuli are given in the appendix.

3.3.7: Expt. 4: Vernier acuity for Gaussian bars flanked by an edge: Stimuli

Patterns were composed of two Gaussian bars designated the target bar, with one flanking cumulative Gaussian edge. The bars were of positive contrast polarity; the edge was either positive contrast polarity (its peak value above the background luminance in the rest of the display) or negative contrast polarity. The flanking edge was 50% of the contrast of the target bars. Extensive pilot trials indicated that this flank contrast did not noticeably increase the vernier threshold, but that a higher contrast of flanking feature was necessary to induce a perceptual shift than in the preceding experiments. The stimulus configuration differs from previous experiments in that there was only one flanking feature. This flanking edge could occur beside either the top or bottom target bar, but always on the right, as seen by the observers. Luminance profiles for the stimuli used are given in the appendix at the end of this chapter.

3.3.8: Expts. 3 and 5: Vernier acuity for Gaussian edges flanked by a bar or an edge: Stimuli

Patterns were composed of two cumulative Gaussian edges, designated the target, with one flanking Gaussian bar. The edges were of positive contrast polarity; the bar was either positive or negative contrast polarity. The flanking bar was 50% of the contrast of the target edges. Vernier targets for edges flanked by edges consisted of two cumulative Gaussian edges as a vernier target, with a cumulative Gaussian edge flank. The target edges were of positive contrast, and the flank edge was either negative or positive contrast. Flank contrast was 50% of the target contrast. Luminance profiles are given in the appendix.

3.3.9: Experiment 6: Interference with interval width perception

The standard stimuli consisted of two Gaussian bars, with space constants of 1.0 min arc, with the peaks of their intensity profiles being separated by 16 min arc. The comparison stimulus was flanked by either a positive or a negative contrast pair of flanking bars. These flank bars were Gaussians, with a space constant of 1.0 min arc, but a contrast of 50% that of the target bars. The subject had to indicate in which interval the wider of the two stimuli occurred.

3.3.10:Viewing conditions

The display was viewed from a distance of 228 cm. Observers were seated with their heads supported by a chin rest, temple and forehead restraints. Viewing was binocular. For the direction of displacement discrimination, the stimuli appeared as two notional 'frames', each with a duration of 500 msec, with the display raster shifted between the two extremely rapidly. The first vernier experiment (expt. 2) had exposure durations of 2000 msec; the following experiments had exposure durations of 500 msec.

3.4: Results

3.4.1:Experiment 1: Direction of displacement discrimination

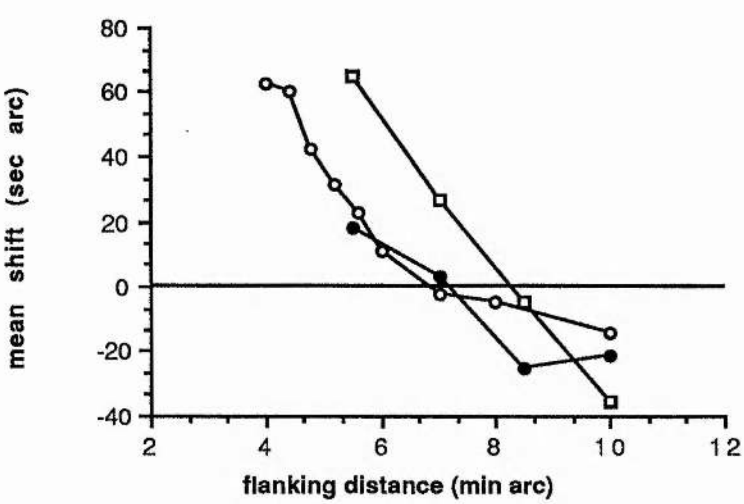

Figure 3.5: stimuli were Gaussian bars: space constant 1.0 min arc, direction of displacement discrimination task: mean shifts obtained for three subjects. The graphs shows the effect of varying the separation of the flank bar from the target bar on direction of displacement discrimination; the data points represent the physical displacement of the target bar which gives a percept of no apparent motion of the target bar when the side of the flanking bar is exchanged at the instant of the displacement (ie. the physical shift of the target bar to null the apparent motion induced by the change in side of the flanking bar). Positive values on the ordinate indicate perceptual "attraction", negative values "repulsion". Contrast of the flanking bars was 25% that of the target bars. Data is shown for three subjects (filled circles: SM, open circles: IRP, open squares: LD). Data points are based on 240 observations.

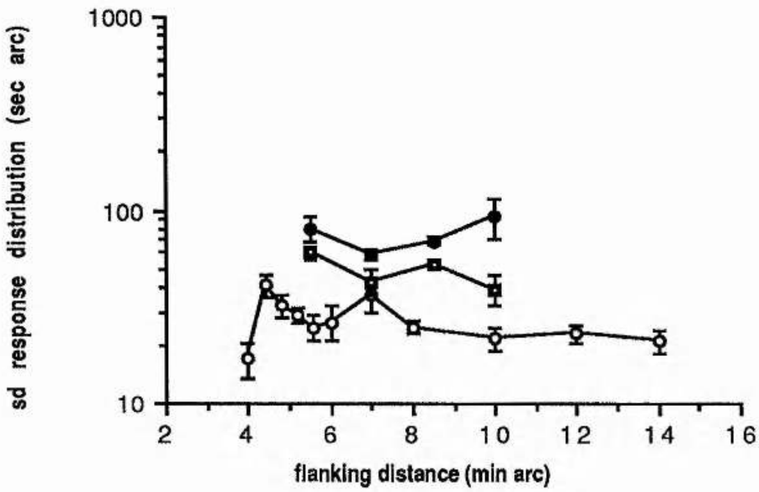

Figure 3.6: Gaussian bars, space constant 1.0 min arc: direction of displacement discrimination thresholds shown for the three subjects (filled circles: SM, open circles: IRP, open squares: LD). Data points are based on 240 observations. Thresholds are represented as the standard deviation of the error response function. There is some indication of an increase in direction of displacement thresholds at flanking distances of less than 5.5 min arc for subject IRP, of roughly the same magnitude as eg. Westheimer and Hauske (1975) where a 3-fold elevation in threshold was seen for flanking distances in the range 2.5-5 min arc for a vernier task. Error bars represent one standard error of the mean; these statistics are computed by the probit analysis program.

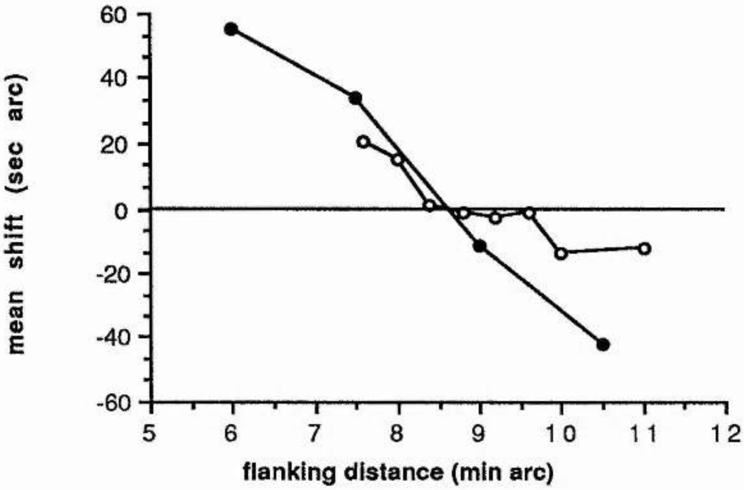

Figure 3.7: Gaussian bars: space constant 2.0 min arc, direction of displacement discrimination task: mean shifts obtained , for subjects IRP (open circles) and SM (filled circles).

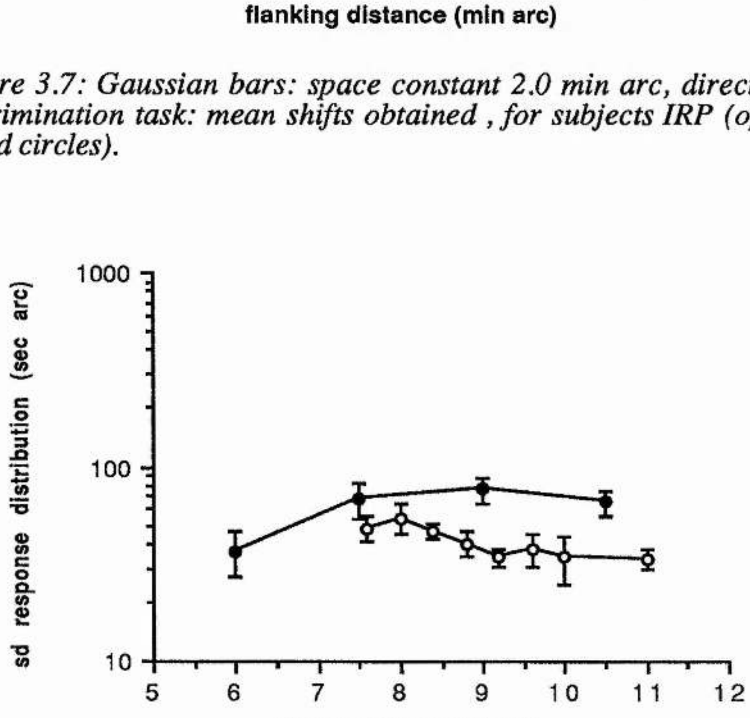

Figure 3.8: Gaussian bars, space constant 2.0 min arc: direction of displacement discrimination thresholds shown for two subjects (filled circles: SM, open circles: IRP). Data points are based on 240 observations. Again, there is some indication of an increase in direction of displacement thresholds at flanking distances of less than about 8 min arc for subject IRP.

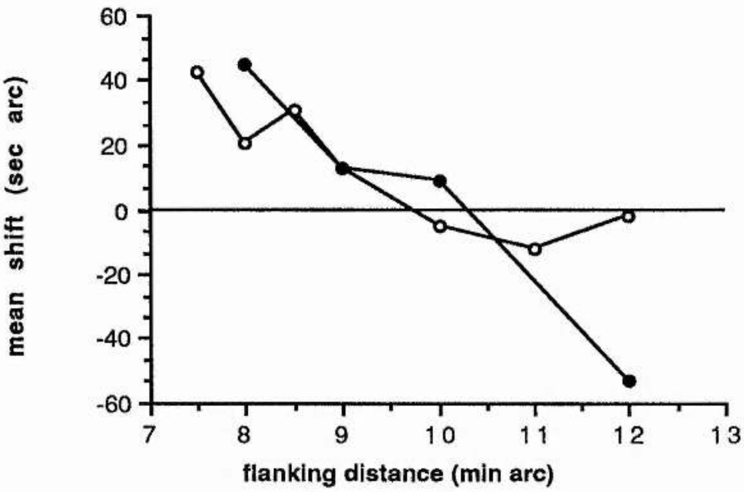

Figure 3.9: Gaussian bars: space constant 2.5 min arc, direction of displacement discrimination task: mean shifts obtained for subjects IRP (open circles) and SM (filled circles).

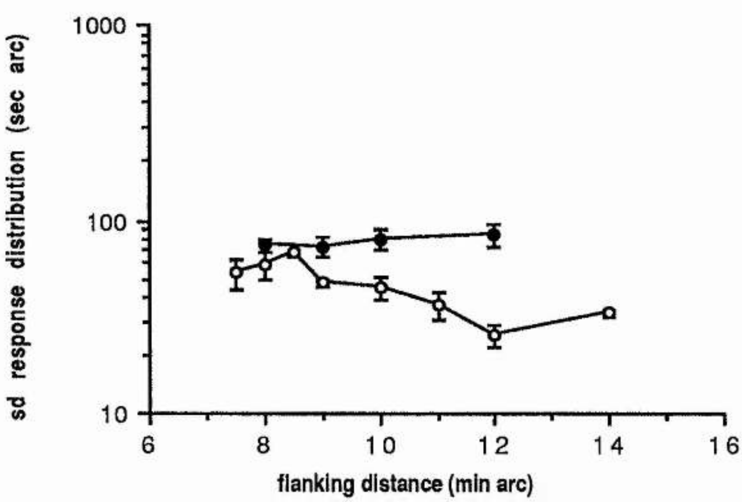

Figure 3.10: Gaussian bars, space constant 2.5 min arc: direction of displacement discrimination thresholds shown for two subjects (filled circles: SM, open circles: IRP). Data points are based on 240 observations. There is some evidence of an increase in thresholds at flanking distances of less than about 12 min arc for subject IRP.

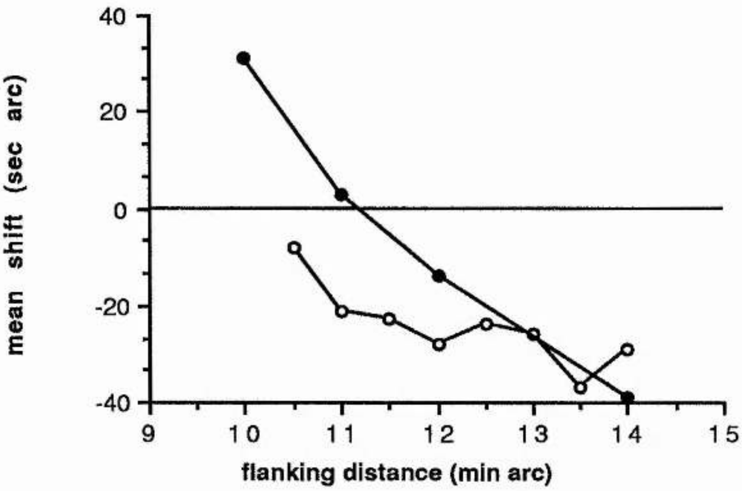

Figure 3.11: Gaussian bars: space constant 3.0 min arc, direction of displacement discrimination task: mean shifts obtained for subjects IRP (open circles) and SM (filled circles).

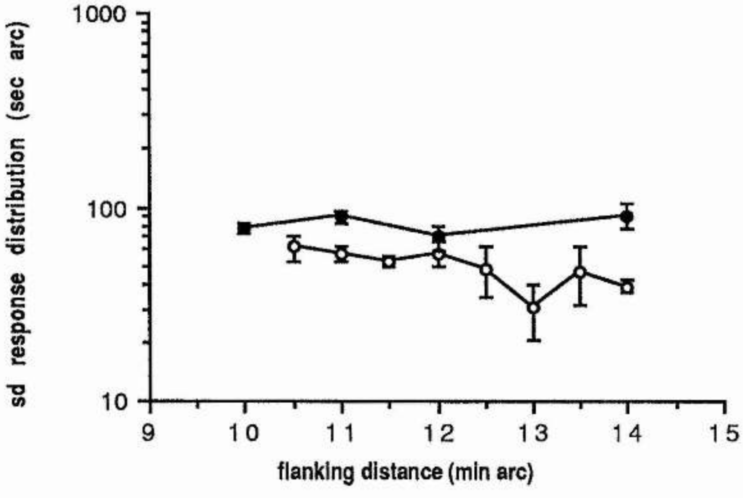

Figure 3.12: Gaussian bars, space constant 3.0 min arc: direction of displacement discrimination thresholds shown for two subjects (filled circles: SM, open circles: IRP). Data points are based on 240 observations. There is some evidence of an increase in thresholds at flanking distances of less than about 12 min arc for subject IRP.

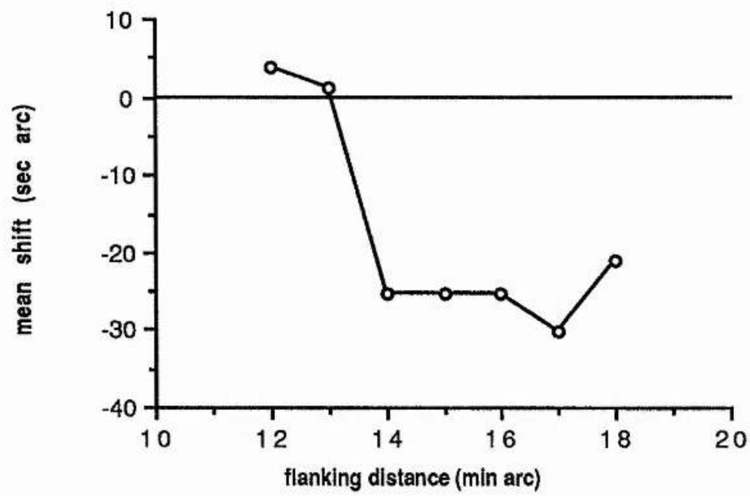

Figure 3.13: Gaussian bars: space constant 4.0 min arc, direction of displacement discrimination task: mean shifts obtained for subject IRP

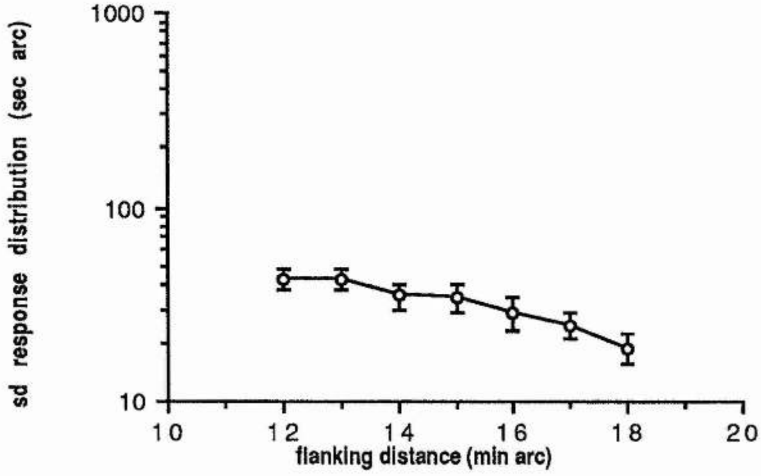

Figure 3.14: Gaussian bars, space constant 4.0 min arc: direction of displacement discrimination thresholds shown for one subject (fIRP). Data points are based on 240 observations.

Mean shifts and direction of displacement thresholds obtained are shown above in figures 3.5-3.14 as a function of the separation of the flanking bar and target bar. Results from all observers show reasonably good correspondence, although observers LM and SD were considerably less experienced on psychophysical tasks than IRP, and consquently their thresholds are generally quite a lot higher. The mean-shift value refers to the displacement of the target bar which gave a percept of 'no-motion', obtained by measuring the spatial thresholds for direction of displacement discrimination rather than a direct nulling technique. Mean-shifts indicate the combined influence of two flank bars on opposite sides of the target bars in the 500 msec time windows before and after the instant of displacement of the target bar. A positive value on the ordinate indicates an attraction of the target towards the flanking feature. The pattern of mean shifts indicate that smaller separations (below about 10 min arc) of the flank and target lead to a perceptual attraction of the target and flank, ie. a biasing of the perceived location of the target bar in the flank direction. At greater separations of target and flank, a perceptual repulsion is observed. This replicates the results of Badcock and Westheimer (1985a). However, the present results depart from those reported by Badcock and Westheimer in several respects, which illustrate previously unreported features of the spatial interference phenomenon; there is some trend towards attraction at greater flanking distances as the blur of the target bar is increased (eg. a 25% width in the extent of the attraction is seen for change in blur space constants from 1.0 to 3.0 min arc). Extent of the attraction zone is not linearly related to the stimulus blur, unlike the observed relation between interference with vernier acuity and eccentricity (Levi, Klein & Aitsebaomo, 1985). Attraction effects were not reliably found when the blur of bars exceeded about 3.0 min arc. There was some indication that the presence of the flanking bar elevated thresholds for this task maximally at at flanking distance of about 8 min arc, evident in the data of IRP. This compares with a figure of 3-5 min arc for vernier acuity for fine lines (Westheimer & Hauske, 1975).

3.4.2: Experiment 2: Vernier acuity for Gaussian bars flanked by Gaussian bars

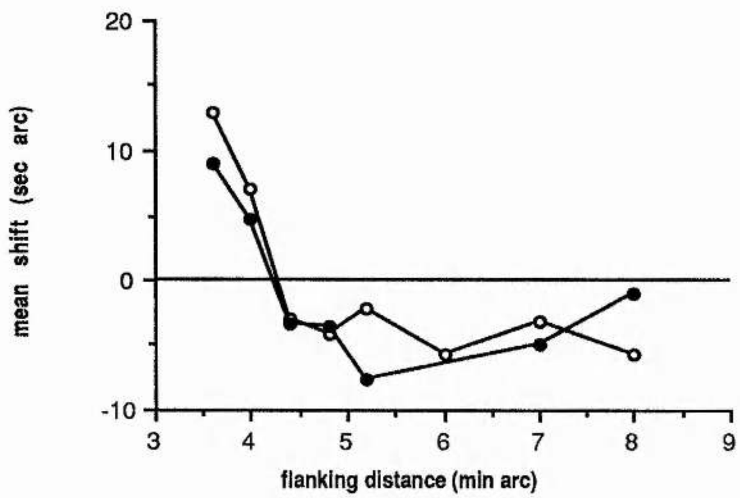

Figure 3.15: Shifts in the point of subjective alignment of two bars in a vernier target (location biases) induced by the presence of flanking bars are presented as a function of the retinal distance between the target and flanking bars. Positive values on the ordinate represent biasing of the target bars in the direction of the flank ("attraction"), negative values a biasing of the target bars away from the direction of the flank ("repulsion"). Data is shown for two subjects; IRP (open circles) and GLC (filled symbols). Both target and flanking bars have Gaussian contrast profiles, with a space constant of 1.0 min arc.

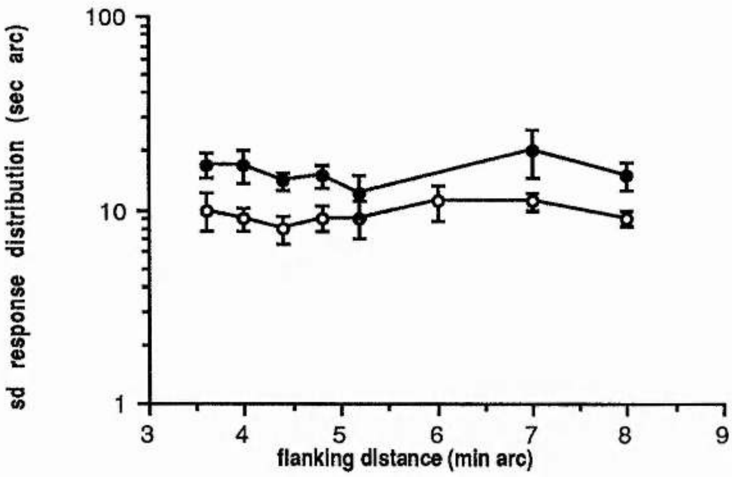

Figure 3.16: Gaussian bars, space constant 1.0 min arc: vernier thresholds shown for two subjects (filled circles: GLC, open circles: IRP). There is no evidence of any increase in vernier thresholds for the smaller flanking distances. Hence the mean shifts obtained in the graph above represent changes in the perceived relative location of the target bars without accompanying changes in the precision of localisation of the target bars.

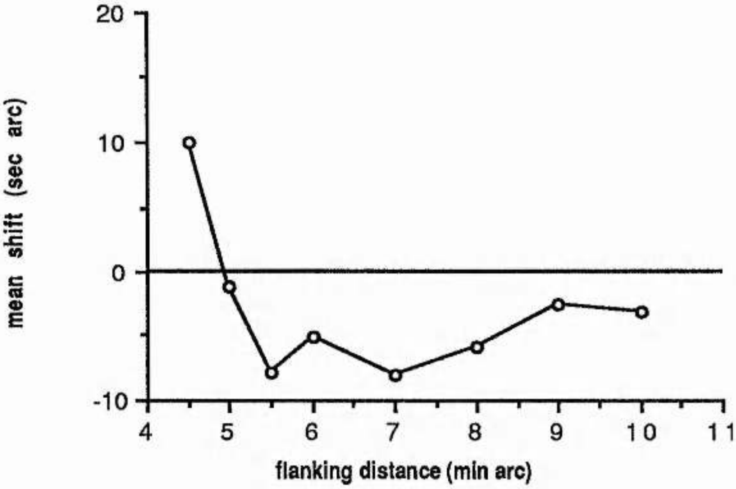

Figure 3.17: Location biases for one subject (IRP) for Gaussian bars with space constant of 1.5 min arc in the presence of flanking bars. Attraction is indicated by a positive value on the ordinate. Repulsion is observed at all but the smallest of flanking distances. At flanking distances smaller than those tested, it becomes increasingly more difficult to resolve the two bars.

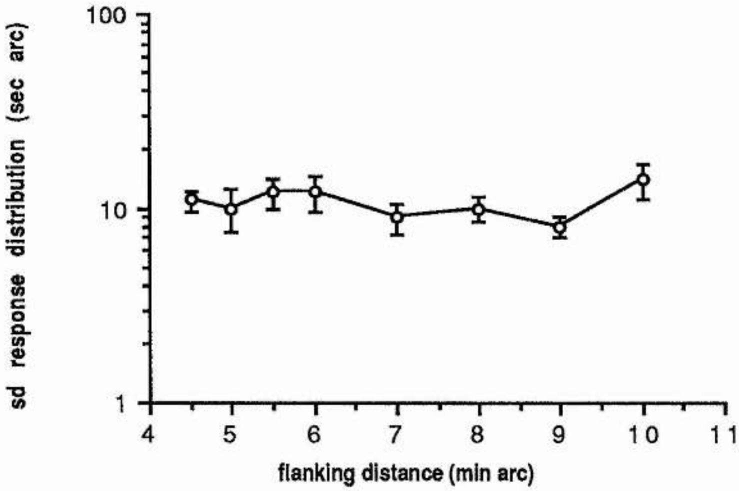

Figure 3.18: Vernier thresholds for localising Gaussian bars with space constant of 1.5 min arc in the presence of flanking Gaussian bars are shown for subject IRP.

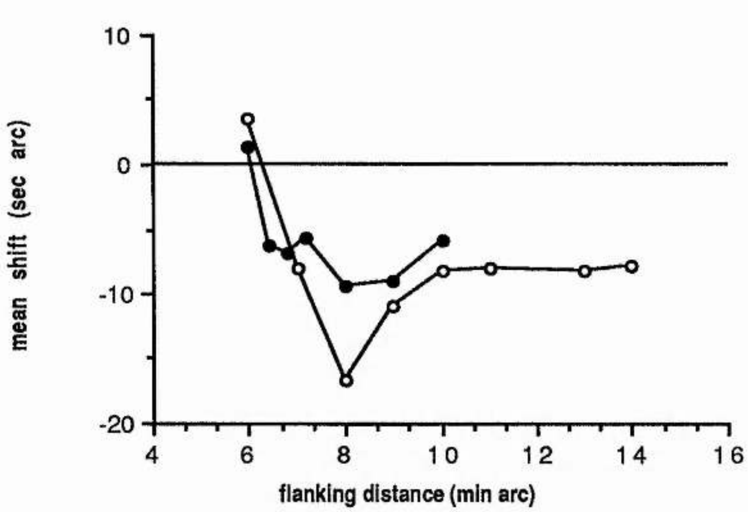

Figure 3.19: Flanked Gaussian bars with space constant 2.0 min arc; location biases are shown for two subjects (IRP; open circles, and GLC; filled circles). Attraction is only weakly present for both subjects at the smallest flanking distances tested.

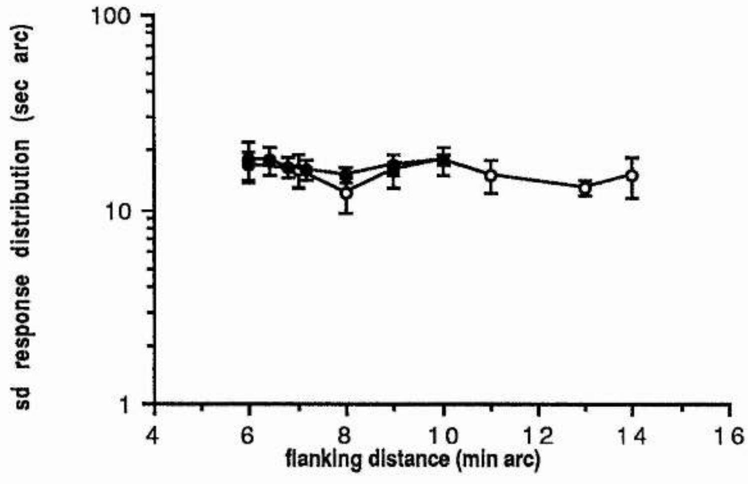

Figure 3.20: Vernier thresholds are shown for flanked Gaussian bars with space constants of 2.0 min arc. Data from two subjects is shown (IRP; open circles, and GLC; filled circles). There is no clear effect of flanking distance on precision of localisation in this data.

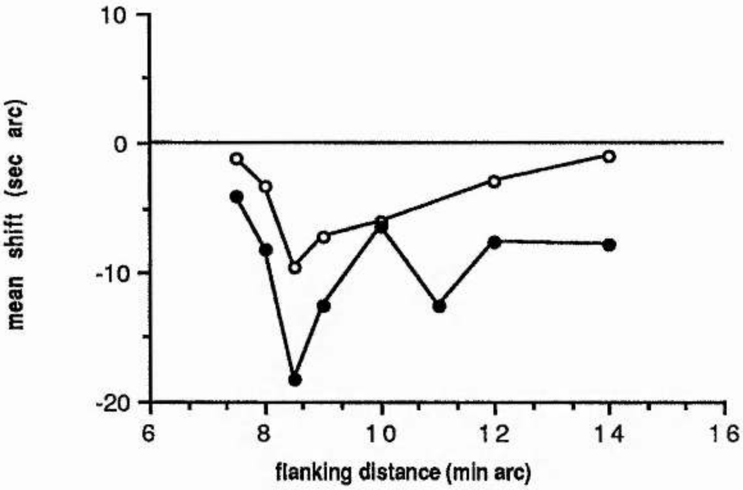

Figure 3.21: Gaussian bars with space constant of 2.5 min arc; location biases are shown for two subjects (IRP; open circles, and GLC; filled circles). Attraction is represented by a positive value on the ordinate, and is not evident at all for either subjects at any of the flanking distances tested.

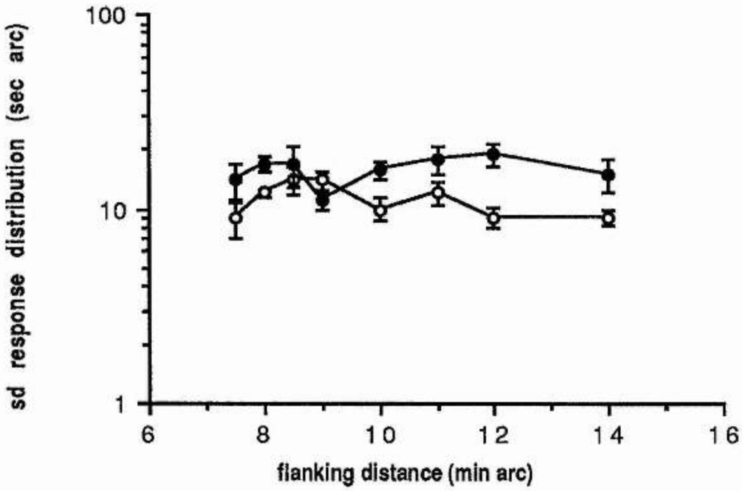

Figure 3.22: Vernier thresholds are shown above for localisation of flanked Gaussian bars with space constants of 2.5 min arc. Data is shown for two subjects (IRP; open circles, and GLC; filled circles). There is some suggestion of elevated thresholds at the smallest flanking distances tested, but this was not evident in the data from the smaller blurs, and is probably illusory and due to noise in the data.

Mean shift values and vernier thresholds obtained are shown in figures 3.15-3.22 as a function of the separation of the flanking bar and the target bar. Results for the observers are closely similar. Note that as in the previous experiment, the perceived shifts determined represent the combined effects of two flanking bars on opposite sides of the upper and lower target bars. These mean shifts were obtained without any noticeable trends in the vernier thresholds for the different flanking distances.

The results of this experiment confirm and extend those of Badcock and Westheimer (1985a). Attraction between resolved bars (bars which could be distinctly perceived as spatially separated objects: the stimulus intensity distribution of the two bars should possess a sufficiently large 'dip' between the luminance maxima for the bars to be resolved, of the order of 10% of the maximum amplitude of the luminance excursion) disappears if the bar is blurred beyond 2-2.5 min arc space constant, and is obtained at flanking distances of less than 6-7 min arc. This was verified for subject IRP and another observer with bars of 4.0 min arc; with bars of this blur, if they could be resolved, then attraction could not be observed.

In the 1st experiment involving a direction of displacement discrimination task, attraction was obtained at greater flanking distances, and with slightly more blurred bars than in the vernier acuity task. This indicates that the mean size of spatial filters mediating direction of displacement discrimination (ie. perception of apparent motion) may be somewhat larger than those subserving vernier acuity.

As a control condition, vernier thresholds and means were collected at different exposure durations for one subject, to discover the effect of exposure duration; since it did not appear to affect the basic pattern of results, this was not pursued further. These experiments were conducted prior to publication of Watt's (1987) results indicating coarse-to-fine processing in early vision, and it did not appear at the time theoretically interesting to pursue these results further, although in retrospect it may have been interesting to have done so, since they conflict with the basic notion of a coarse-to-fine analysis of spatial relations. The data is reported here for completeness in figures 3.23-3.26, but since it is based on only data from the author (IRP), only tentative conclusions should be drawn.

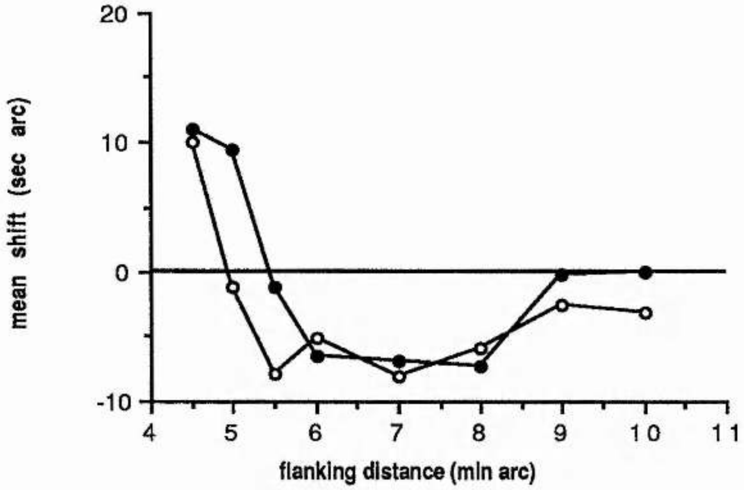

Figure 3.23: Mean shifts obtained under two exposure durations from subject IRP; filled circles are for a duration of 100 msec, and open circles are for a duration of 2000 msec. There is a remarkable correspondence between the two sets of data. Space constants of the target and flanking bars was 1.5 min arc.

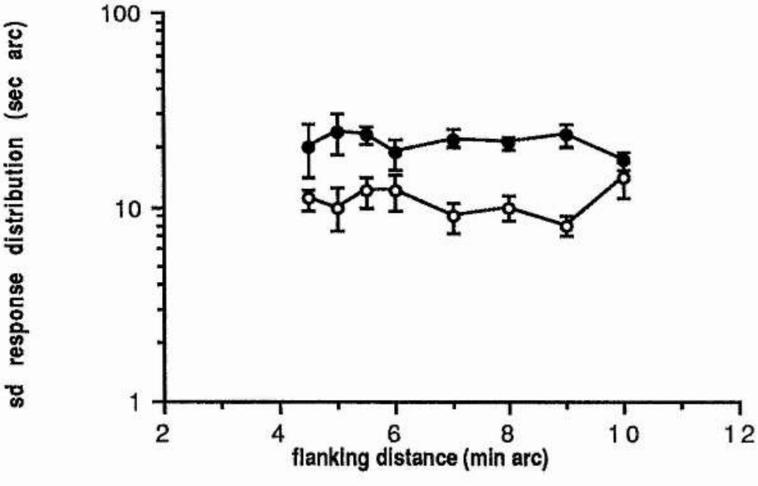

Figure 3.24: Localisation thresholds for two different exposure durations for subject IRP. Open circles are for the 2 sec exposure duration, filled circles represent data from the 100 msec condition. The space constant for the bars was 1.5 min arc. Thresholds were somewhat elevated by shortening the exposure duration, but neither function has any obvious structure to it.

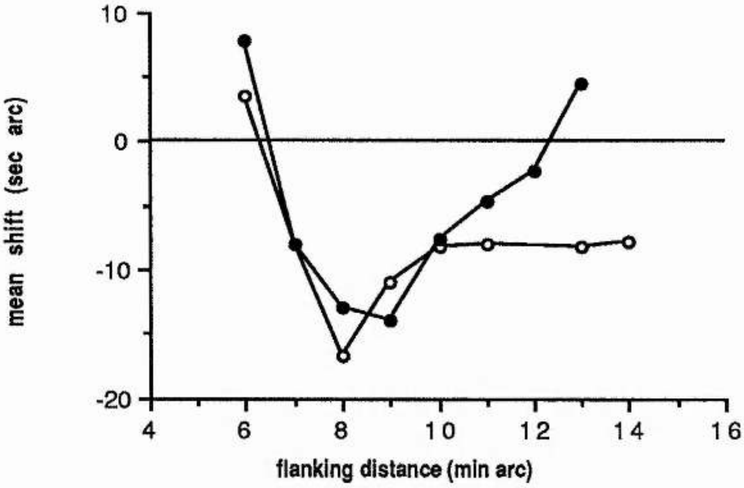

Figure 3.25: Mean shifts from subject IRP, for two different exposure durations. Both target and flanking bars had Gaussian contrast profiles with space constants of 2.0 min arc. Open circles are for the 2 sec exposure duration, filled circles for the 60 msec exposure duration .

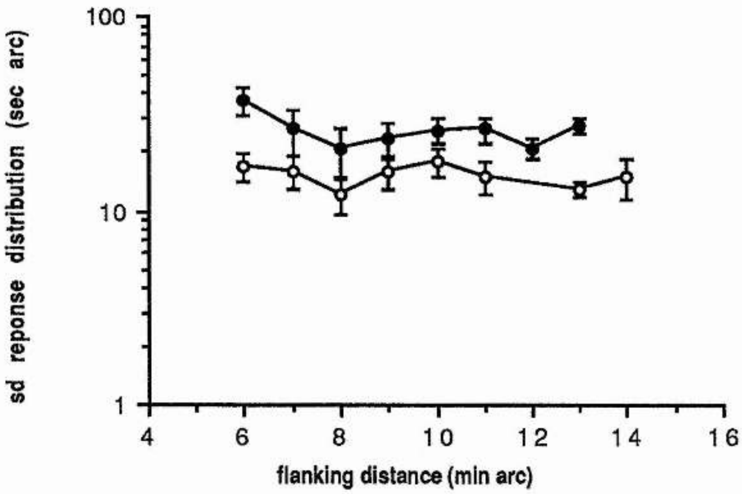

Figure 3.26: Vernier thresholds from subject IRP, for two different exposure durations. Both target and flanking bars had Gaussian contrast profiles with space constants of 2.0 min arc. Open circles are for the 1 sec exposure duration, filled circles for the 60 msec exposure duration condition. Note that there is some indication of decreased precision of localisation for the 60 msec exposure duration condition at smaller flanking distances.

3.4.3: Experiment 3: Vernier acuity for edges flanked by bars

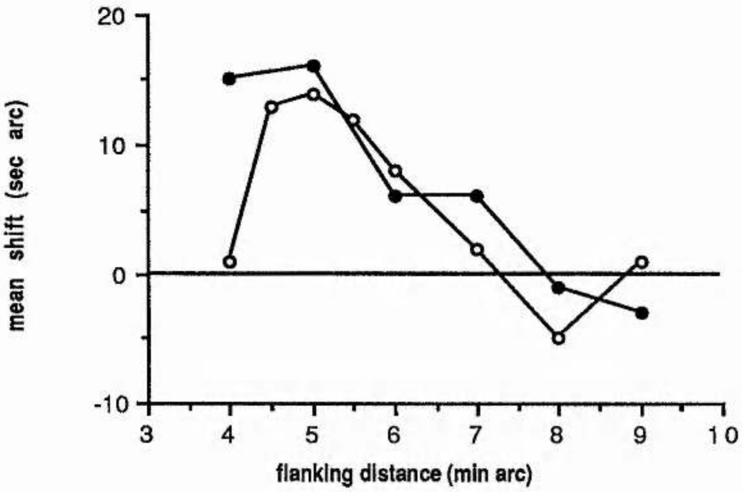

Figure 3.27: Mean shifts are shown for edge targets, with a single Gaussian bar flank, which varied between flanking the upper and lower edge. Data is shown for two subjects (IRP; open circles and GLC; filled circles). The flanking bar had a contrast 50% that of the edge and had a negative polarity- i.e. the appearance of the target was of a dark bar beside a bright edge. Attraction is evident at small flanking distances for both subjects despite the polarity of the target feature, which is the reverse of the results when the target is a bar (eg. Badcock & Westheimer, 1985a,b).

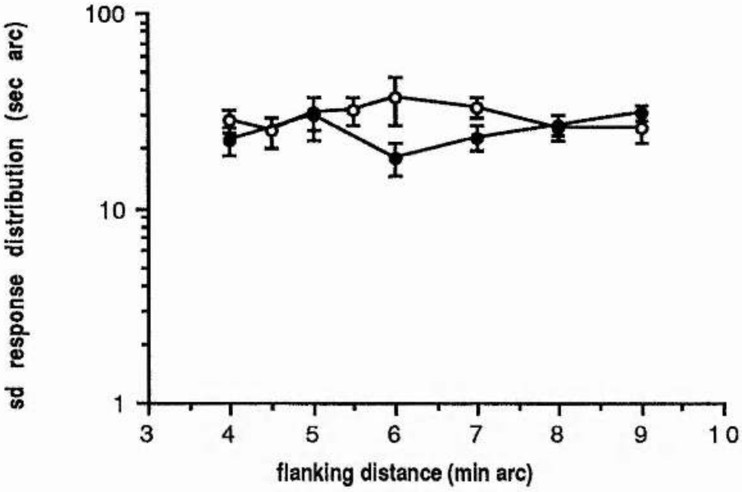

Figure 3.28: Localisation thresholds are shown for edge targets, with a single Gaussian bar flank, which has a negative contrast polarity (and is 50% the contrast of the edge). Data is shown for two subjects (IRP; open circles, and GLC; filled circles). There is no clear dependency of precision of localisation on the separation of the target edges and the flanking bar. The edges and bar had space constants of 1.0 min arc.

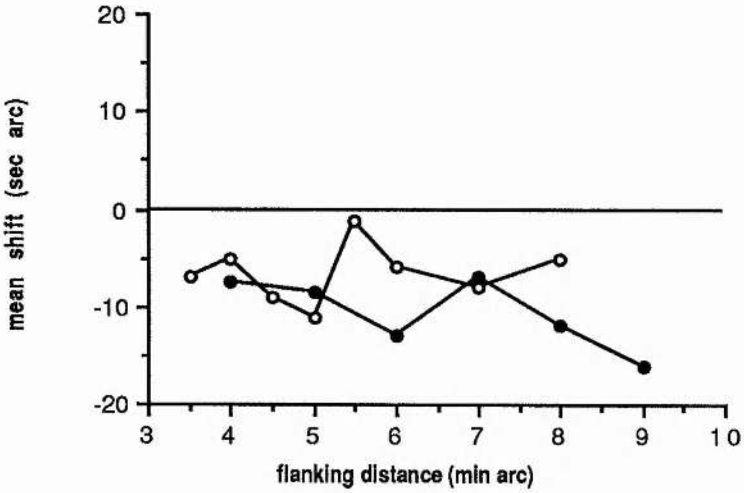

Figure 3.29: Mean shifts are shown for edge targets, with a single Gaussian bar flank. Data is shown for two subjects (IRP; open circles and GLC; filled circles). The flanking bar had a contrast 50% that of the edge and had a positive polarity- of contrast.

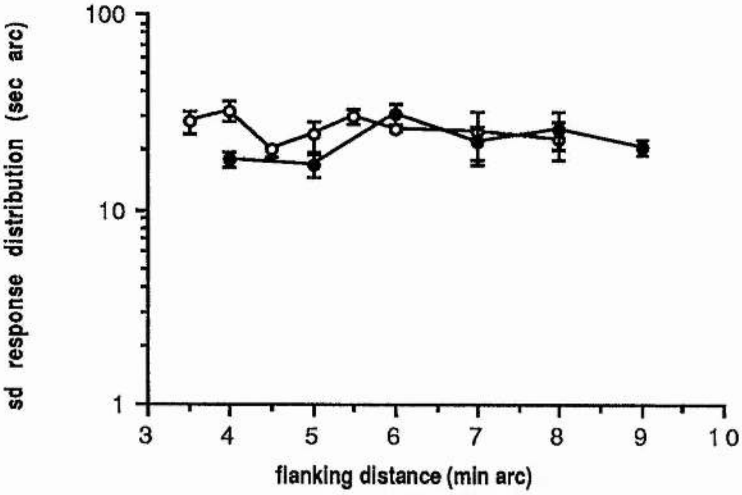

Figure 3.30: Localisation thresholds are shown for edge targets, with a single Gaussian bar flank with positive contrast polarity (50% the contrast of the edge). Data is shown for two subjects (IRP; open circles, and GLC; filled circles). Again, there is no clear dependency of precision of localisation on the separation of the target edges and the flanking bar. The edges and bar had space constants of 1.0 min arc.

Mean shifts values and thresholds obtained are displayed in figures 3.27-3.30 as a function of the flanking distance between the target cumulative Gaussian edge and flanking Gaussian bar. The results are unlike those presented for line-line interactions (Badcock & Westheimer, 1985a,b) or in the previous experiments, for either the bar-bar or bar-edge interactions. It is evident that the effects of polarity of contrast of the flanking feature are reversed if the feature to be localised is an edge. Badcock and Westheimer demonstrated only repulsion of lines if the flank and target are opposite in contrast, but in the present experiment, attraction is seen between a positive contrast blurred edge and a negative contrast blurred bar. A positive contrast blurred edge is repelled by a positive contrast blurred bar.

This indicates that the interaction between bars and edges is asymmetric: in the next experiment, it was demonstrated that a positive contrast target bar was attracted towards a positive contrast flanking edge. The present experiment demonstrates that this must be accompanied by a simultaneous shift in the perceived location of the edge away from the target bar in these circumstances.

In comparison with the data from a bar flanked by an edge, an edge which is to be localised is distorted in position more by an adjacent bar, than a bar to be localised which is flanked by an edge. The shifts in perceived location were reliably larger when the feature to be localised was an edge, rather than a bar; partly, this may reflect the higher thresholds for edge localisation (Watt & Morgan, 1983a), although subject's phenomenal reports suggested that the edge localisation was genuinely impaired more by a flanking feature than bar localisation.

3.4.4: Experiment 4: Vernier acuity for Gaussian bars with edge flanks

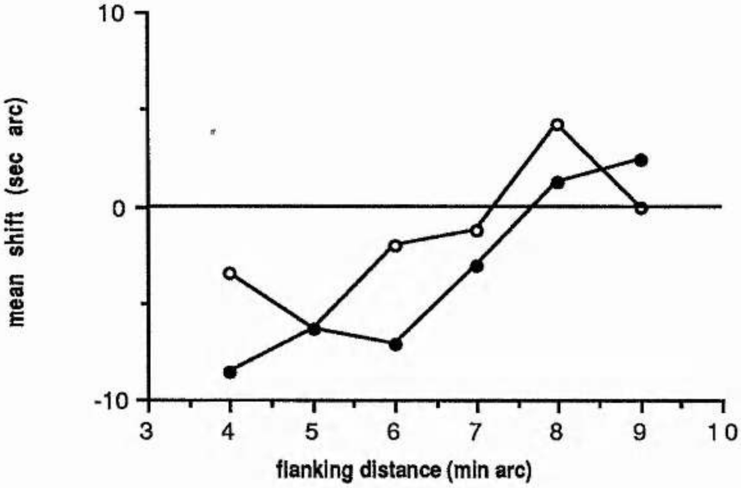

Figure 3.31: Location biases are shown for two observers for positive contrast Gaussian bar targets flanked by a negative contrast edge. Space constant of the bars and edge was 1.0 min arc. Data is shown for IRP (open circles) and MEW (filled circles). The subjects indicate a consistent pattern of bias, with the bars being biased away trom the flanking edge at least at the smaller flanking distances tested. This is similar to the data of Badcock and Westheimer (1985a), where they showed that a positive contrast target line is repelled by a negative contrast flanking line.

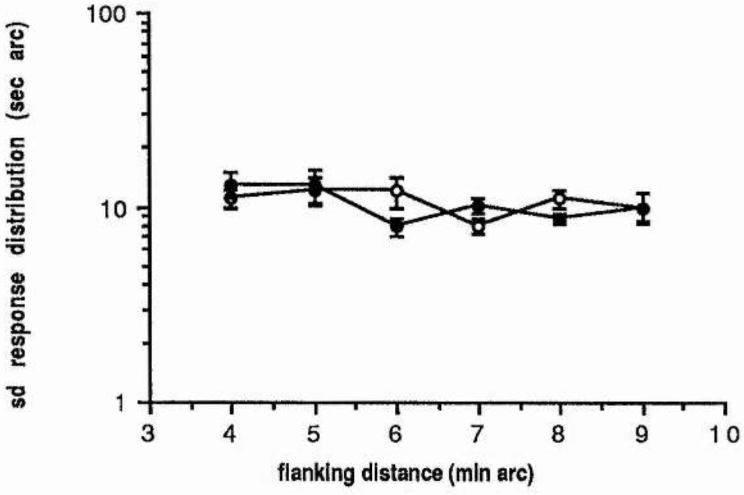

Figure 3.32: Vernier thresholds are shown for Gaussian bars flanked by an edge, the bars and edge having a space constant of 1.0 min arc. The bars were of positive contrast polarity and the flanking edge of negative contrast polarity. (Open circles: IRP, filled circles: MEW).

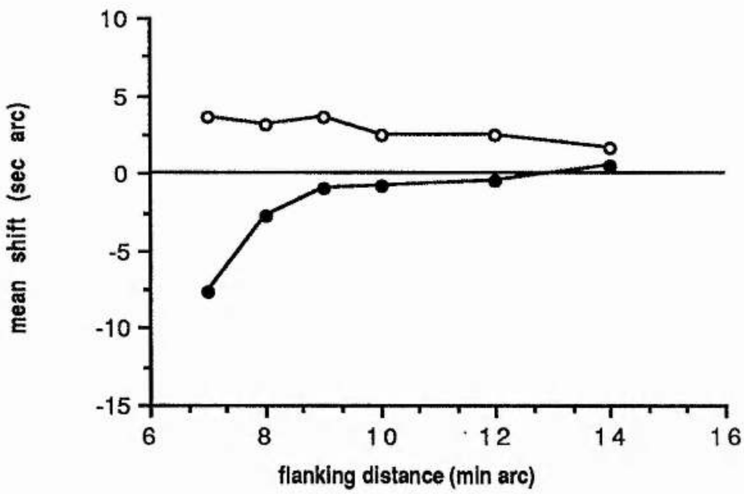

Figure 3.33: Location biases for localising Gausian bars with a single edge flank. Above is shown ***pooled*** *data from two observers, IRP and MEW. The open circles represent data from the condition with a positive contrast polarity of edge flank, and the filled circles the data from the condition with a negative contrast polarity edge. Space constant of the bar targets and edge flank was 2.0 min arc. The data indicate that a positive contrast bar is attracted to an edge of the same contrast, whereas it is repelled by an edge of different contrast.*

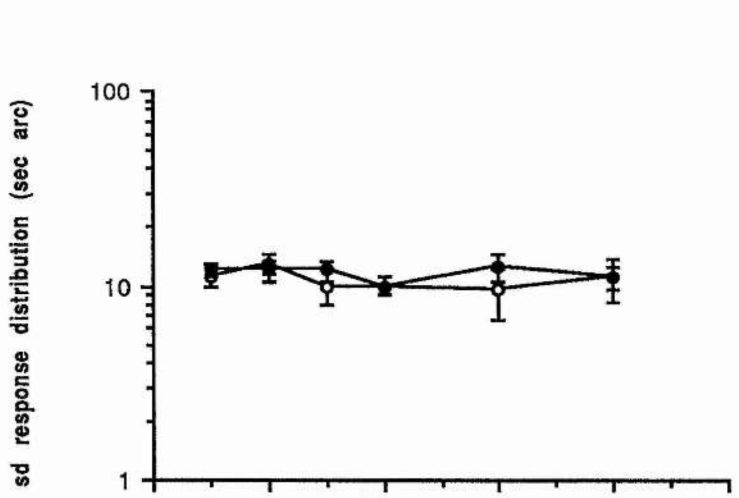

Figure 3.34: Vernier thresholds for 1.0 min arc Gaussian bars with a flanking edge with the same space constant are shown pooled across the two subjects (MEW and IRP); the open circles are for the positive contrast flank condtion, and the filled circles for the negative contrast flank condition.

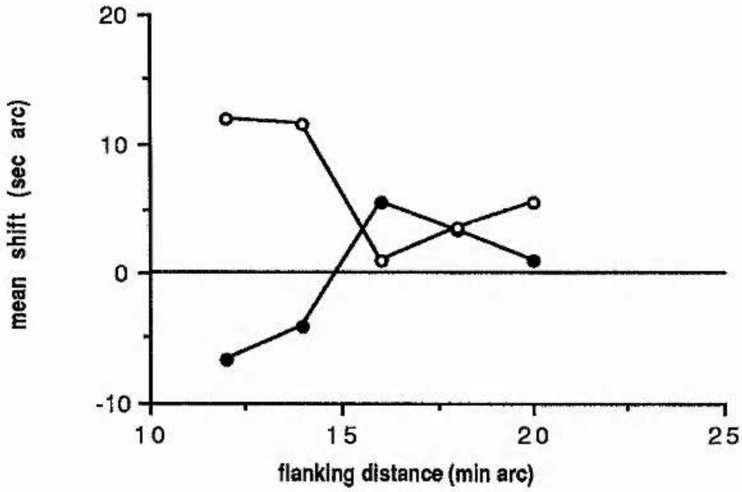

Figure 3.35: The data above is ***pooled*** *across the two observers (IRP and MEW). The open circles are for the condition where the flank is of positive contrast, and the filled circles are for the negative contrast flank condition. The space constant of the bars was 4.0 min arc.*

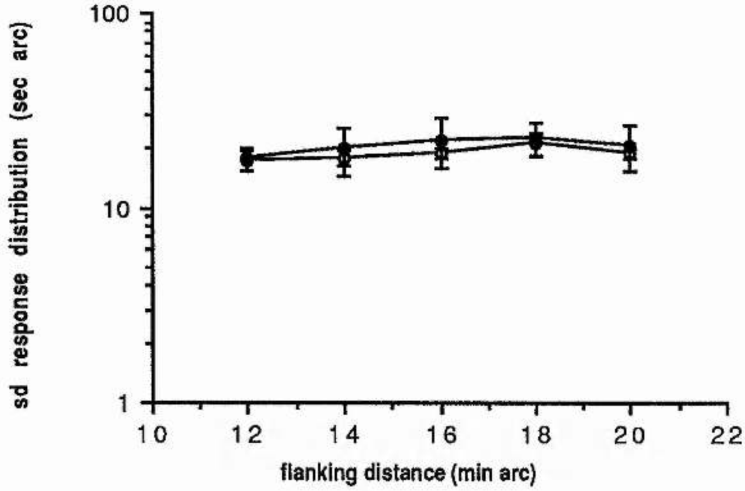

FIgure 3.36: Vernier thresholds for localisation of Gaussian bars are shown pooled across the two subjects. The open circles represent data from the positive contrast flank condition ,and the filled circles data from the negative contrast flank condition. Space constant of the target bars and flanking edge was 4.0 min arc.

Mean shift values obtained are presented in figures 3.31-3.36 as a function of the distance between the centre of the cumulative Gaussian flanking edge and the Gaussian target bar. Results of the observers are similar. The perceptual location of a blurred bar is influenced by the presence of a nearby blurred edge. However, the effects are smaller than those reported in the first two experiments, and also smaller than observed when the target feature is an edge. This is of course partly due to the use of only one flanking feature in this experiment, which would roughly halve the size of the induced shift in location of the target features. Both observers were also of the opinion that it was easy to localise a bright bar in the presence of an adjacent edge, and the small size of the mean-shifts obtained may reflect this phenomenal impression.

Attraction is observed for small separation of the target bars and a bright flanking edge, confirming an earlier study which had broadly similar stimulus configurations (Ganz 1964). Perceptual repulsion of the target bar and flank edge was observed at small separations when the flanking feature is a dark edge. The general pattern is similar to bar-bar interactions, although it appears that repulsion is not present at greater flanking distances for either polarity of flank. The data is not firm enough to support a stronger conclusion.

3.4.5: Experiment 5: spatial interference for edge targets with edge flank

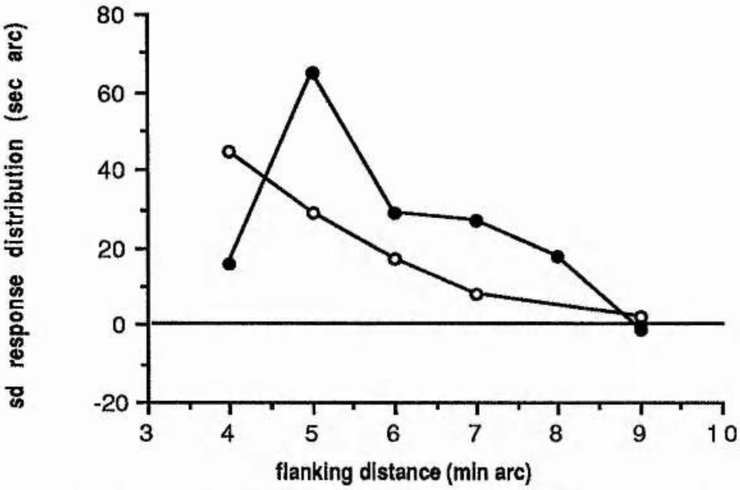

Figure 3.37: Mean shifts are shown for edge targets, with a single cumulative Gaussian edge flank, which varied between flanking the upper and lower edge. Data is shown for two subjects (GLC; open circles and FMRM; filled circles). Space constant of the edges was 1.0 min arc. The flanking edge had a contrast 50% that of the target edges and had a negative polarity- i.e. the appearance of the target was of a dark region in one quadrant, beside two brighter quadrants. Attraction is evident at small flanking distances for both subjects despite the polarity of the target feature, echoing the results where the flank feature was a bar.

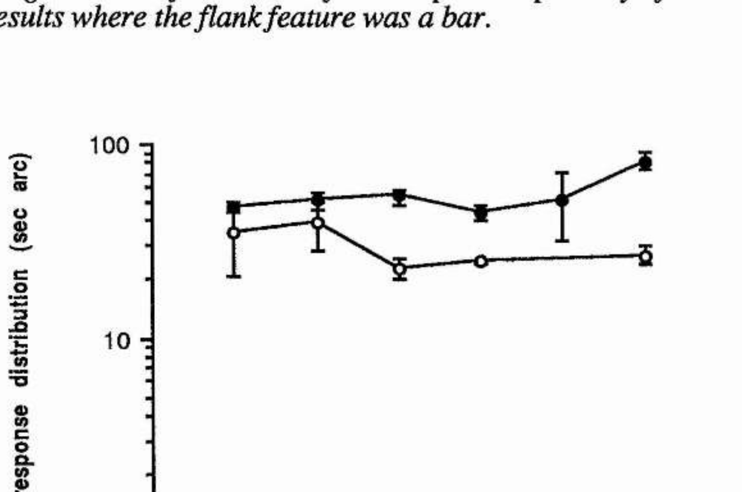

Figure 3.38: Localisation thresholds are shown for edge targets, with an edge flank; the flanking edge had a contrast 50% that of the target edge and a negative polarity. Data is shown for two subjects (GLC; open circles, and FMRM; closed circles). Space constant of the edges was 1.0 min arc.

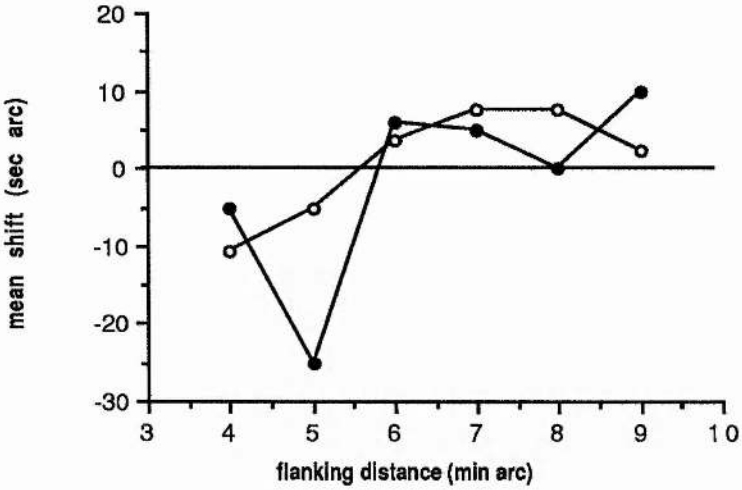

Figure 3.39: Mean shifts are shown for edge targets, with a single cumulative Gaussian edge flank. Data is shown for two subjects (GLC; open circles and FMRM; filled circles). The flanking edge had a contrast 50% that of the target edges and had a positive polarity. Space constant of the edges was 1.0 min arc.

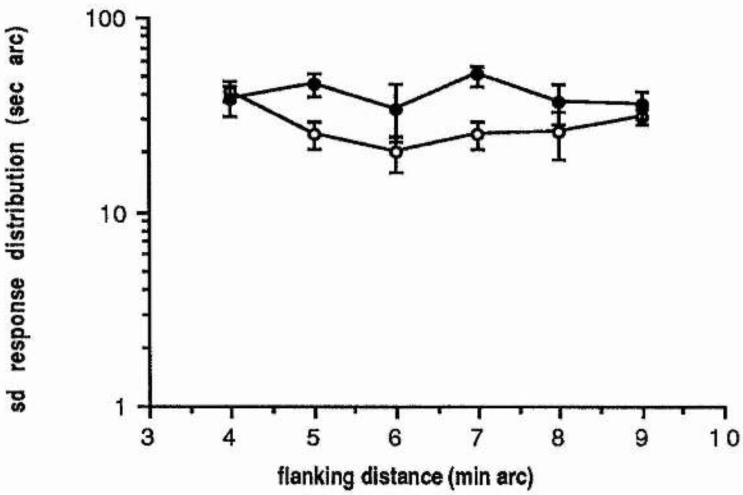

Figure 3.40: Localisation thresholds are shown for edge targets, with an edge flank; the flanking edge had a contrast 50% that of the target edge and a positive polarity. Data is shown for two subjects (GLC; open circles, and FMRM; closed circles). Space constant of the edges was 1.0 min arc.

Mean shift values and thresholds are presented in figures 3.37-3.40 above. The basic result of experiment 3 is confirmed, that a positive contrast flank leads to some repulsion of the edge target at small flank-target separations, and a negative contrast flank leads to attraction, at small separations, in the opposite direction to interference with localisation of bar targets. The data also suggest that blurred edge localisation is more affected by the presence of a flanking features than blurred bar localisation. The measured shifts for the two subjects shown here and a third were much larger than those obtained in experiment 3, with bar target and one edge as flank. There is some indication of increases in thresholds for edge localisation at flanking distances of less than about 6 min arc in the data of subject GLC; subject FMRM, who was relatively inexperienced on psychophysical judgments, does not shown any clear effect of flanking distance on localisation thresholds.

3.4.6: Experiment 6: spatial interference with interval width discrimination

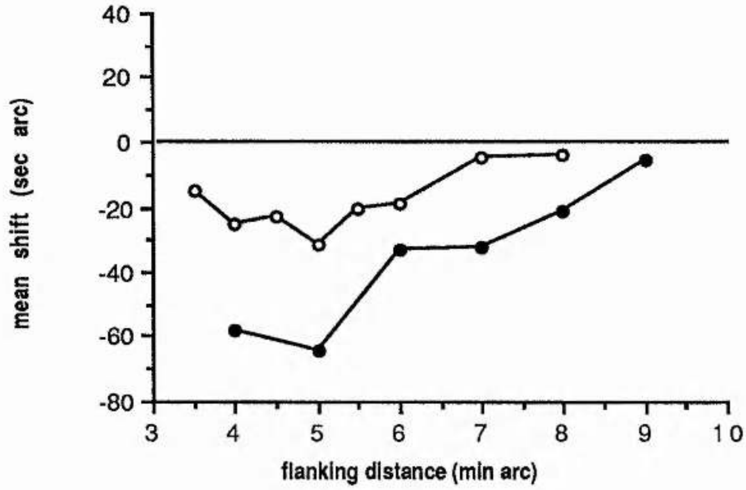

Figure 3.41: Changes in the perceived width of a spatial interval as a function of the separation of the flanking bars are shown above; the flanking bars had a negative contrast polarity and were 50% the contrast of the target bars. All bars were Gaussians with a space constant of 1.0 min arc. The flank bars were placed outside the interval; the reference stimulus for this task was a pair of unflanked bars. A negative value on the ordinate indicates a phenomenal shrinking of the interval; ie. a biasing of the location of the target bars away from the side of the flank bars, equivalent to the the 'repulsion' component observed in vernier tasks. Data is shown for two subjects (IRP; open circles, and MEW; filled circles). Both subjects show a similar effect of the flanking bars, which is to decrease the apparent width of the spatial interval, although the magnitude of the effect differs somewhat for the two observers.

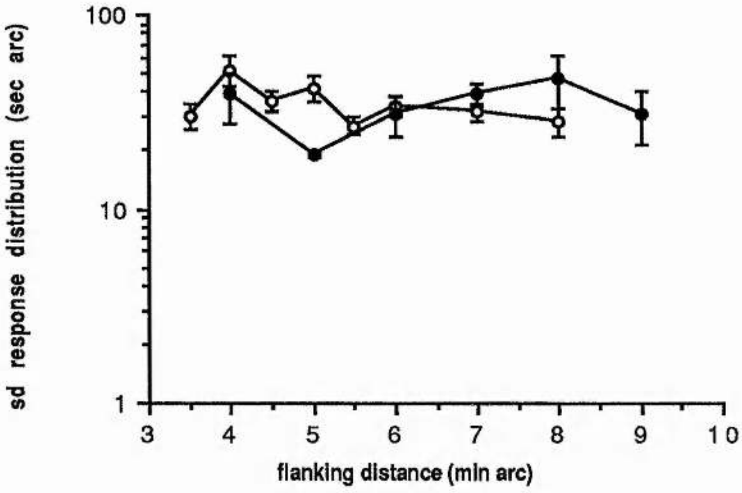

Figure 3.42: Interval discrimination thresholds are shown for Gaussian bars with a space constant of 1.0 min arc, separated by 16 min arc. The flanking bars had a contrast of 50% that of the target bars, and had a negative polarity of contrast. The target bars had positive contrast polarity; thus the stimulus took the appearance of two bright bars flanked on the outside by two dark bars. Data is shown for two subjects (IRP; open circles, and MEW; filled circles).

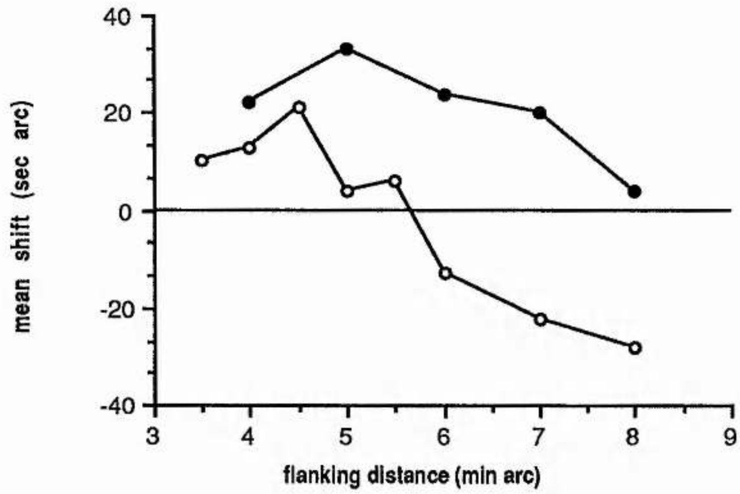

Figure 3.43: Changes in the perceived width of a spatial interval are shown as a function of the flanking distance, for flanking bars with positive contrast polarity. The bars were Gaussians with a space constant of 1.0 min arc. Data is shown for two subjects (IRP; open circles, and MEW; filled circles). There are quite large inter-subject differences, but both show a similar trend, with an increase in the perceived width of the interval at small flanking distances, and the effect of reversing the polarity of the flank contrast is the same for both observers.

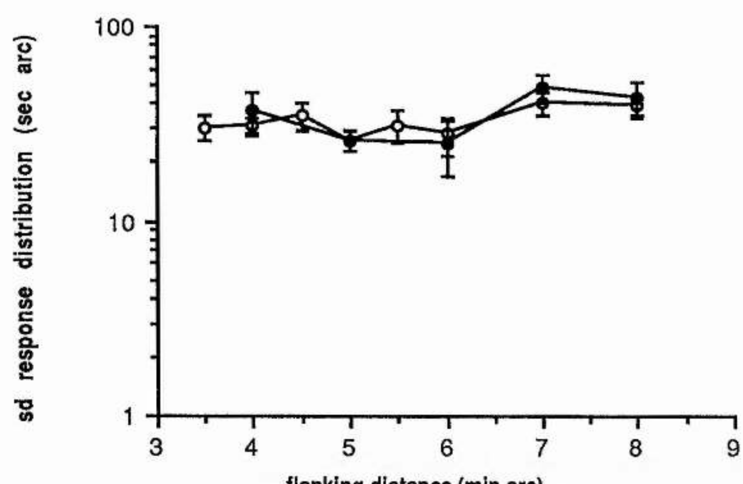

Figure 3.44: Interval discrimination thresholds are shown for Gaussian bars with a space constant of 1.0 min arc, separated by 16 min arc. The flanking bars had a contrast of 50% that of the target bars, and had a positive polarity of contrast. The target bars had positive contrast polarity. Data is shown for two subjects (IRP; open circles, and MEW; filled circles). Both subjects show an indication of elevated thresholds at the greatest flanking distances.

Mean shift values and interval thresholds, corresponding to changes in the perceived width of the flanked interval relative to a unflanked standard, and discrimination accuracy , are presented above in figures 3.41-3.44. A positive value on the ordinate (attraction of target and flank) indicates that the interval was perceived as being wider than the unflanked reference interval, since the flank bars were on the outside of the spatial interval. The pattern of interactions obtained are similar to those found in interference with vernier acuity, confirming the result of Levi and Westheimer (1987), who used two line targets with an inner perturbing line, unlike the four line targets presented here. Levi and Westheimer used only a positive contrast internal flank; the data from this experiment indicate that a negative contrast flank has the same effect as in a vernier acuity situation, ie. inducing repulsion at small flank-target separations. Since fairly large shifts in mean were obtained in the present experiment, it is surprising that the data of Morgan and Ward (1985) did not demonstrate any effect of having randomly positioned flanks parallel to the lines delimiting the target interval in one of their experiments; however, the data from some of the experiments in chapter 5 gives an indication as to why their result might have been expected. The data on precision of interval discrimination from this experiment show no clear trends, although there is some indication of a decrease in precision with the larger flanking distances when the target and flank bars were of the same contrast polarity.

3.5: Discussion

The data from the preceding experiments indicate that it is possible to reliably influence perceived direction of apparent motion and relative location or interval width, without affecting the precision with which these tasks can be performed, replicating some previous results (eg. Badcock & Westheimer, 1985a,b; Levi & Westheimer, 1987). An interesting difference is between the effect of blur on spatial interference with displacement discrimination and vernier acuity. With the displacement discrimination task, attraction was evident at greater distances and with slightly more blurred bars than in vernier acuity. This would be expected if the displacement task is mediated by larger scale filters than the vernier task; ie. if the abrupt change in the spatial position in the displacement task is registered by mechanisms with transient temporal characteristics; the retinal distance over which flanking bars enhance the detectability of a target bar in subthreshold summation experiments increases when the temporal modulation is transient (Wilson, 1978).

The attraction zone does not scale with stimulus blur, indicating that the effect is a function of retinal distance rather than a consequence of stimulus geometry. From the interference with vernier acuity for bars, the attraction effect between resolved features cannot be observed beyond quite small blurs of bar (about 2.0 min arc). If the target and flank could not be resolved, then the flank would be expected to induce a perceived shift in the same direction as the attraction between resolved features by creating a bar with an asymmetric contrast profile. The relative location of this bar would be somewhere near the centroid of the retinal intensity distribution (by analogy with the data of Watt, Morgan & Ward, 1983a, where the perceived relative location of a bar with a skewed orthoaxial contrast profile was measured), but it seems helpful to make a distinction between this and attraction between resolved features. The contribution of larger filters which cannot resolve the flank and target bars might contribute to the observed location biases, if the larger filters contribute positional estimates of the target bar. On the other hand, Watt (1984) suggested that spatial interference could be explained by distortions in the location of zero-crossings in the output of a second-differential filter to closely spaced stimulus features. The possibility that the perceptual effects might be generated by other spatial primitives (centroids and extrema) was also investigated. The location of these features in the response of a Difference of Gaussians filter (approximating a Laplacian of a Gaussian) to the stimulus was determined. Additionally, a one-dimensional form of MIRAGE

(Watt and Morgan, 1985) was implemented, to compare its predictions with the data.

3.5.1: Model predictions

One possible way in which the distortion in perceptual location observed in spatial interference could occur is a shift in the position of features in the output of a filter small enough to resolve the target and flank features, but for which there is a distortion in the output of the local filter bank caused by the presence of the nearby features. This supposes that the system has independent access to the outputs of various scales of filter. The first model tested assumes that the shifts in perceived location of intensity changes result from distortions in primitive representations derived from the output of a single scale of filter. However, this model fails dismally, regardless of the choice of spatial primitive. It is not possible to explain why attraction should be seen if only one scale of filter is employed; a repulsion effect is generated as soon as the two bars in the target can be resolved. Despite the wide variety of models tested, there are none which can explain why the attraction zone should be seen if only one scale of filter is involved in relative localisation of intensity changes. The suggestion of Toet et al (1988) that their data indicates only one spatial scale of filter is involved in detecting and representing intensity changes cannot account for the present result.

An alternative to the above is to assume that the visual system does not have independent access to the outputs of filters of different scales, but only to some inflexible combination of their outputs (Hess, Pointer & Watt, 1989; Watt, 1988; Watt & Morgan, 1985, 1984). Using the filter scale parameters suggested by Watt and Morgan (1984), predictions were generated about the perceived relative location of the bars. For closely spaced bars, it is not clear how the output of the MIRAGE operation should be interpreted. The feature which was chosen for pairs of bars was the extremum in noise- bounded regions of response. Note that for closely spaced bars, there are two extrema in one of the noise-bounded response distributions, which are very close to the veridical centres of the bars. Taking the centroid of such a response distribution seems to be discarding the useful information from these extrema (which explain the present results very accurately). Watt (1988) suggested taking a weighted mean of the location of centroids of different zero-bounded regions of MIRAGE's response distribution. This worked well enough when there were two regions of positive response (for instance, in the situation where the system is localising an edge) but it was not apparent how to implement this when there was only one 'fused' region of positive response.

3.5.3: Predictions: Vernier acuity and direction of displacement discrimination

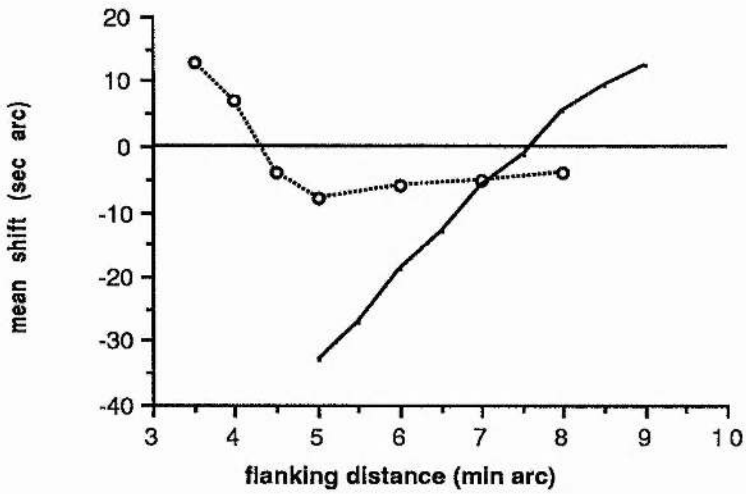

Figure 3.45: The graph above shows data from experiment 2, involving flanked vernier bars. The bars had space constants of 1.0 min arc; the open circles are pooled data from the two subjects, and the continuous line is the predictions of a model based on the zero-crossings in the response of a filter with a space constant of 1.25 min arc; the bars are localised midway between the zero-crossings in the filter response. This model predicts that repulsion effects would be observed at smaller flanking distances than attraction effects; this basic pattern is seen for all sizes of filters which can resolve the two bars.

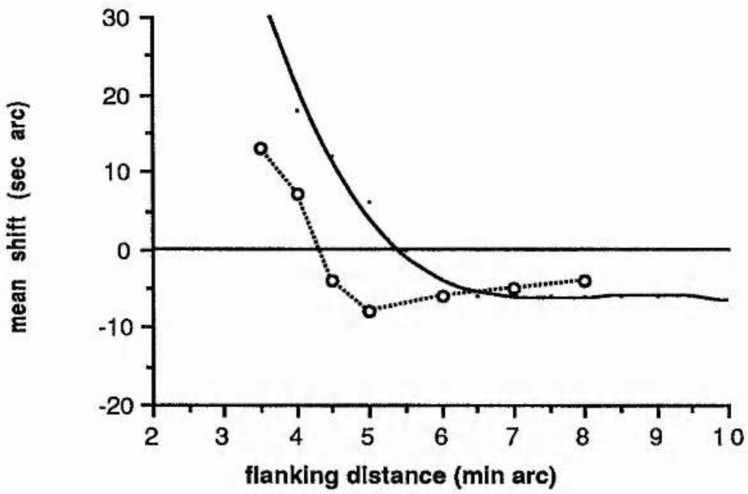

Figure 3.46: Data from the second experiment, involving localisation of flanked Gaussian bars. The bars had space constants of 1.0 min arc; pooled data is represented by the open circles. The continuous line represents the prediction of a MIRAGE model, in which the location of the bars is represented by the extrema in the output of the MIRAGE algorithm, with filter space constants from 0.35-2.8 min arc (after Watt & Morgan, 1984). The model correctly predicts attraction at small flanking distances and repulsion at greater flanking distances, and the transition between these regions is close to that determined empirically.

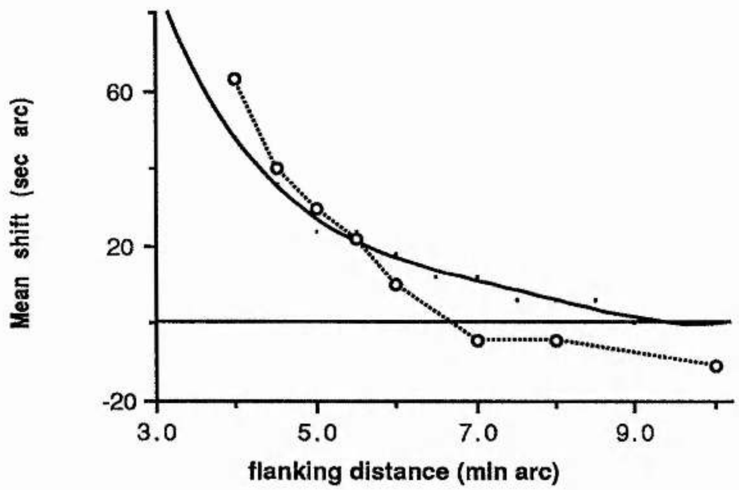

Figure 3.47: Data from experiment 1, involving direction of displacement discrimination for Gaussian bars. Predictions of a MIRAGE-like model, assuming that the visual system localises extrema in the output of the MIRAGE algorithm. Filters are Laplacian of Gaussians, with space constants from 0.7- 5.6 min arc. Good qualitative agreement is shown with the data, which is for Gaussian bars with space constants of 1.0 min arc.

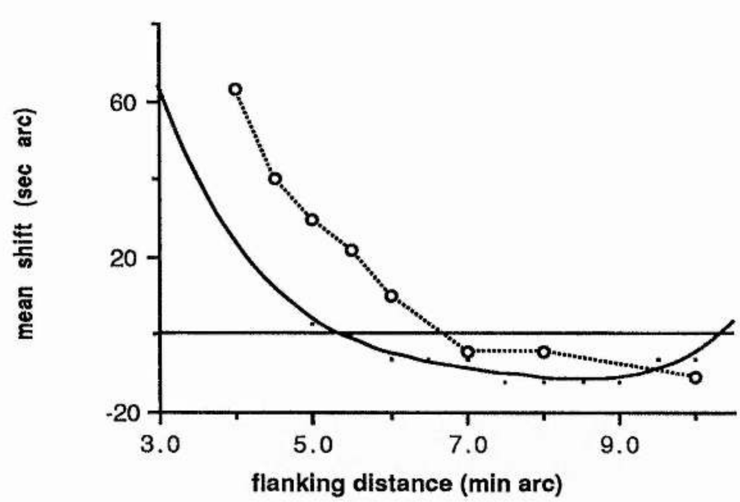

Figure 3.48: Same data set; predictions of MIRAGE when the largest filter has a space constant of 2.8 min arc; it is apparent that a larger filter must be involved to give a function which crosses the x-axis at the empirically observed flanking distance, somewhat larger than 2.8 min arc but perhaps a little smaller than 5.6 min arc.

Experiment 5: localisation of edges in the presence of an edge flank

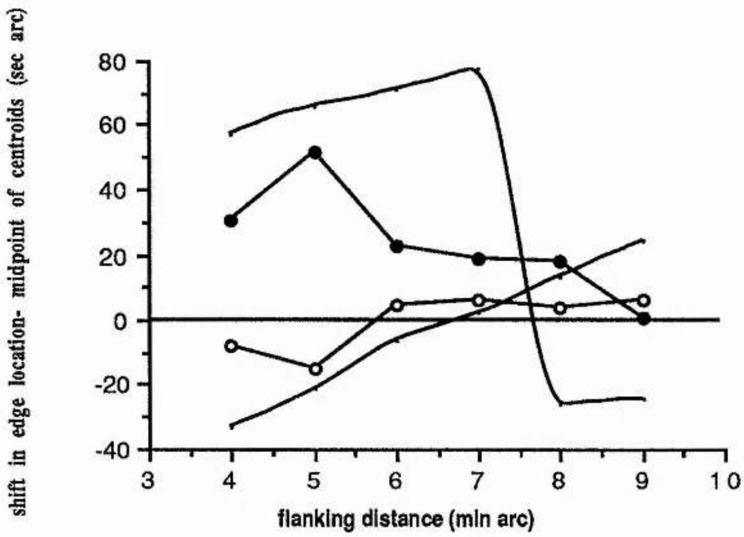

Figure 3.49: The above graph shows the results of simulating the effects of the MIRAGE operation on the location of an edge, in the presence of a flanking edge. Both the target edge and the flank edge have space constants of 1.0 min arc. The filter space constants chosen were 0.35-2.8 min arc, in octave steps. Edge location was determined as the midpoint of the line joining the locii of centroids of opposite sense in the output of the MIRAGE operation. The circles on the graph represent pooled data from two naive subjects. The open circles are for the positive contrast flank, and the filled circles are for the negative contrast flank. The uppermost function is the simulated data for an edge target with a negative contrast flank, the lowermost function the simulated data with a positive contrast flank.

The simulated data reproduces the major qualitative aspects of the empirical data; the data from the negative contrast flank displays the correct attraction effect (a positive value on the ordinate) at the smallest flanking distances, with some indication of a transition to repulsion at greater flanking distances. The simulated data from the condition where the flank is of positive contrast have the same overall pattern as the empirical data- repulsion at the smallest flanking distances, and some indication of attraction at greater flanking distances.

Predictions are shown in figures 3.3.45-3.49. MIRAGE explains several aspects of the data. It explains successfully the data from Experiments 1 and 2; the transition from attraction to repulsion occurs at the appropriate flanking distance. Note that the parameters used by Watt and Morgan (1984) to fit edge localisation data in a completely different experiment, (ie. filter space constants from 0.35-2.8 min arc) succeed here in accounting for the observed perceived shifts in relative location in experiment 2 of this chapter. MIRAGE predicts correctly that attraction between resolved bars should not be observed if the stimulus is sufficiently blurred.

The difference between the results of the direction of displacement discrimination task and the vernier acuity task can be understood as a shift in scale of the filters involved in MIRAGE. This is supported indirectly by the observation of Watt and Morgan (1984) that an additional filter with space constant of 5.6 min arc was needed to fit the data when one of the subjects allowed his eyes to move freely over the display. There is also much evidence from detection experiments for a scale-shift when temporal frequency is increased (eg. Wilson, 1978; Wilson & Bergen, 1979). Modelling of the results of the present displacement experiment suggest that an octave shift downwards in the peak frequencies of all of the filters involved in MIRAGE is adequate to explain the direction of displacement discrimination data. A range of filters from 0.7-5.6 min arc explains the displacement data well, with a range from 0.35-2.8 min arc giving a good account of the vernier acuity data from Experiment 1.

MIRAGE explains the data from localisation of edges, in which the effects of the contrast of the flank were inverted relative to localisation of bar targets. For this modelling, the centroids were chosen as spatial primitives, in line with Watt and Morgan's (1985, 1984) suggestion. Combining the information across different scales, and defining edge location as the midpoint of two centroids of opposite sense (after Watt & Morgan, 1985) explains the main qualitative aspects of the data; mean edge location is biased in the direction of the negative contrast flanking feature, which will manifest as perceptual attraction of dark flanks for bright edge targets, which was observed in experiments 3 and 5 with 4 different subjects. Addtionally, the MIRAGE centroids model predicts a transition from repulsion to attraction for positive contrast edge flanks (the reverse of the situation when the target is a bar)- this was exactly what was observed empirically. Note that other models (eg. localising the zero-crossing in a single scale of filters) cannot explain this pattern of results at all.

3.6: Summary

Data was collected from spatial interference with direction of displacement discrimination for blurred bars, interval width discrimination, and vernier acuity for blurred bars and edges, using bar and edge flanks. The effect of blurring the target and flank bars is not to scale the width of the attraction zone. Rather, the attraction zone remains a fairly constant retinal distance across, suggesting that the processing underlying spatial interference reflects a fixed method of localising intensity changes which depends on retinal distance and is not scale-invariant.

Direction of displacement discrimination and vernier acuity display some difference in the extent of the attraction zone and the effects of blur on the pattern of spatial interactions obtained. For vernier acuity, attraction effects disappeared for resolved bars blurred beyond space constants of 2.5 min arc, but attraction could still be obtained with bars of this blur in a displacement discrimination task, and at 50% greater flanking distances.

Spatial interference with localisation of blurred edges displays a novel effect of the contrast of the flanking feature; for small separations of target and flank, the edge to be localised is attracted towards a dark flank feature (either a bar or an edge), and the edge is repelled by a bright feature. Spatial interference with interval width perception for intervals delimited by bars is similar in some respects to interference with vernier acuity for blurred edges (already noted by Levi & Westheimer, 1987) but confirmed for flanks outside the spatial interval and for dark flanks.

The effects of blur, and the spatial interactions between edges and bars, are correctly predicted by MIRAGE. Data was modelled using similar parameters to Watt and Morgan (1984), ie. filters with space constants from 0.35-2.8 min arc, these having DOG sensitivity profiles and a centre-surround ratio of 1:1.6 (ie. approximately a Laplacian of a Gaussian). To explain the data from the displacement task, the filters' peak frequencies were shifted downwards an octave in spatial frequency.

Appendix: Stimuli from Chapter three

Experiment one: direction of displacement discrimination

Luminance profiles were given by

$$L(x) = L_b[\,1 + c.w(x)\,]$$

where L_b = background luminance,
c = contrast (0.45),

$$w(x) = A\,[\,\exp(\frac{-x^2}{2s^2}) + B\exp(\frac{-(x-\mu)^2}{2s^2})\,],$$

where A = scaling constant, $B = 0.25$,
s = space constant of both bars,
μ = flanking distance

The spatial spread of the Gaussians was truncated at +/- four times their standard deviation, which contains 99.99% of the total energy of the Gaussian distribution. The luminance profiles can be seen overleaf, and a greyscale representation. Note that the greyscale image must be viewed from approximately **5 metres** to give the same size of retinal image as in the experiment (for 1.0 min arc bars). At close viewing distances, the presence of Mach bands interferes with perception of these greyscale patterns. Note also that the stimuli were presented in a much wider field than is shown in these representations. The limited resolution of the Mathematica routine used to display the greyscale images prevented a closer analogue being printed to the images displayed on the screen. The displacement task consisted of two images presented in rapid succession, such that the central bar had a small displacement in one direction or the other; the flanking side was changed in side at the instant of the displacement.

Experiment Two: vernier acuity for Gaussian bars with bar flanks

Luminance profiles for the stimuli were given by:

$$L(x) = L_b[\,1 + c.w(x)\,]$$

where Lb = background luminance,
c = contrast (0.45),

$$w(x) = A\,[\,\exp(\frac{-x^2}{2s^2}) + B\exp(\frac{-(x-\mu)}{2s^2})\,],$$

where A = scaling constant, B=0.25,
s = space constant of both bars,
μ = flanking distance,

for all y values in top half of screen.

For lower half of screen,

$$w(x) = A\,[\,\exp(\frac{-(x-\partial)^2}{2s^2}) + \exp(\frac{-(x+\mu-\partial)^2}{2s^2})\,],$$

A, s, μ as above,
∂ = offset, to left or right.

The Gaussian bars were truncated at +/- four times the space constant of each bar.

Experiment three: vernier acuity for Gaussian bars, flanked by edges

Luminance profiles for the stimuli were given by:

$$L(x) = L_b\,[\,1 + c.w(x)\,]$$

where L_b = background luminance,

c = maximum contrast,

and $$w(x) = A.[\exp(\frac{-(x-\mu_b)^2}{(2s^2)}) + B.\int \exp(\frac{-(x-\mu_e)^2}{(2s^2)})dx]$$

where A,B are constants,

s = space constant of bar and edge,

μ_e = mean position of edge,

μ_b = mean position of bar (=0 for upper bar).

for an unflanked bar, B = 0.

Stimuli from experiment 5 were identical to this but with the B term (B=0.5) in front of the first expression rather than the second.

Experiment four: vernier acuity for edge targets, with bar flanks

Luminance profiles for the stimuli were given by:

$$L(x) = L_b\,[\,1 + c.w(x)\,]$$

where L_b = background luminance,

c = maximum contrast,

and $$w(x) = A.[\int \exp(\frac{-(x-\mu_e)^2}{(2s^2)})dx + B.\exp(\frac{-(x-\mu_b)^2}{(2s^2)})\,]$$

A,B are constants,

s = space constant of bar/edge,

μ_e = mean position of edge,

μ_b = mean position of bar.

for an unflanked edge, B = 0. For a flanked edge, B=0.5.

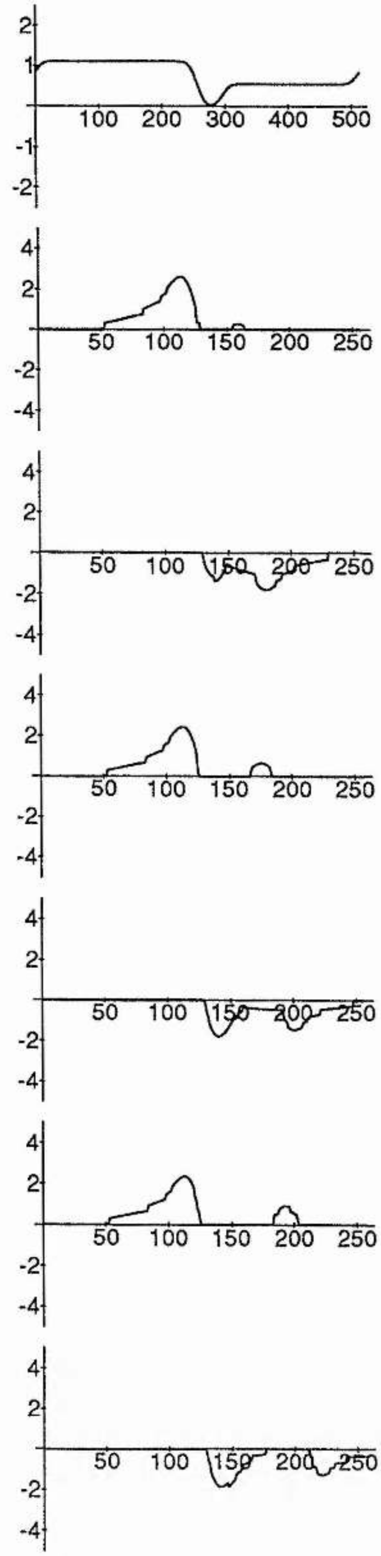

The diagrams above show the output of a MIRAGE operation, when the input stimulus is a target edge with a positive contrast edge flank at 50% the contrast. The stimulus is filtered with LoG filters having space constants from 0.35-2.8 min arc. The responses are normalised and summed, after halfwave rectification. The flanking distance increases from top to bottom (4,6 and 8 min arc), and the positive and negative components of the filter response are shown on separate axes. Centroids are computed for the regions of response, which are truncated at the standard deviation of the response. Edges are located at the midpoint between centroids of opposite sense.

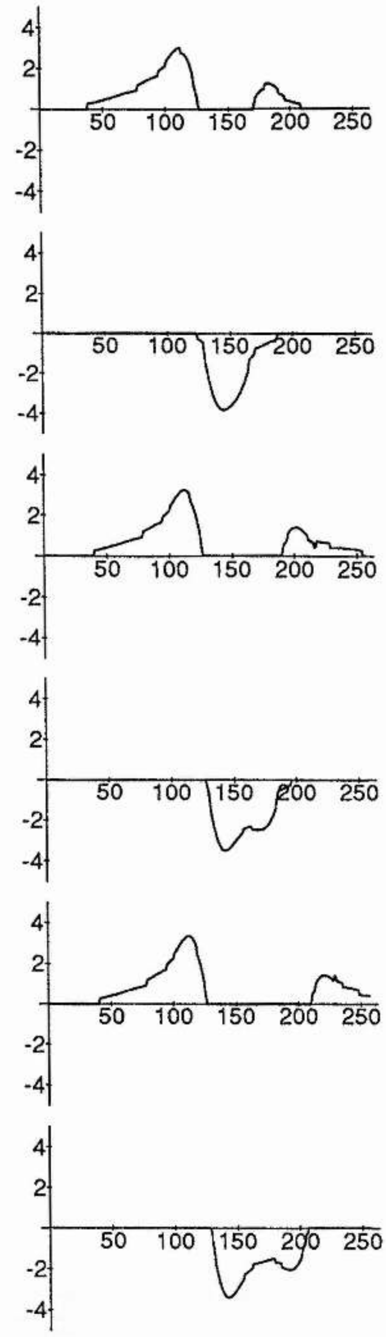

The diagrams above show the output of a MIRAGE operation, when the input stimulus is a target edge with a negative contrast edge flank at 50% the contrast. As in the previous example, the stimulus is filtered with LoG filters having space constants from 0.35-2.8 min arc, the responses are normalised and summed, after halfwave rectification. The flanking distance increases from top to bottom (4,6 and 8 min arc), and the positive and negative components of the filter response are shown on separate axes.

stimuli

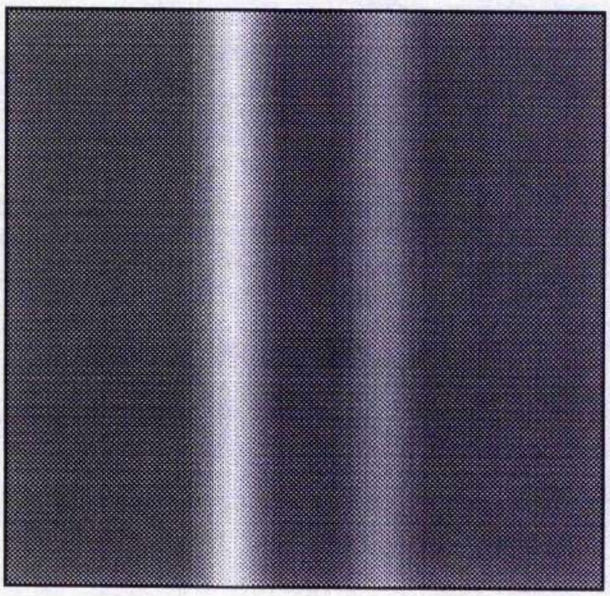

stimuli

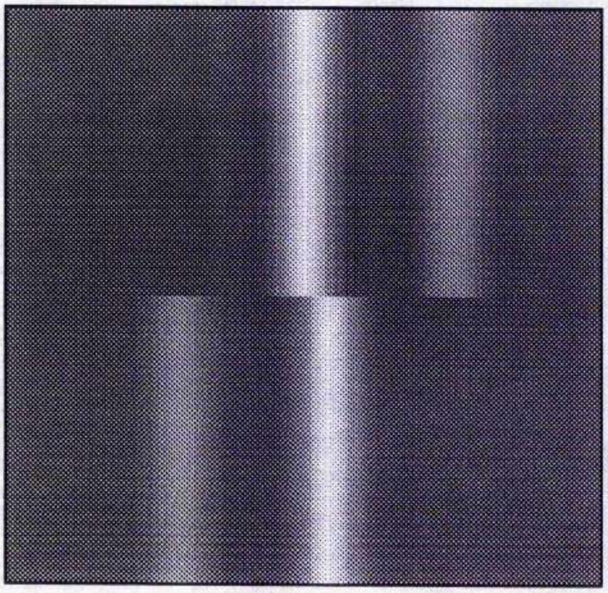

stimuli

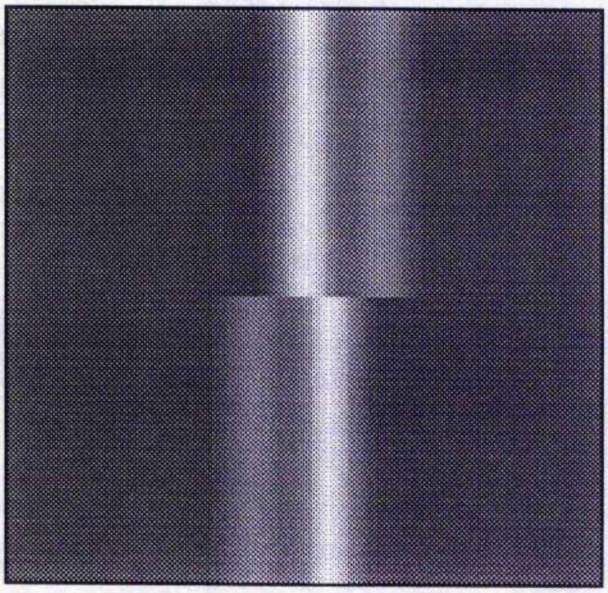

stimuli

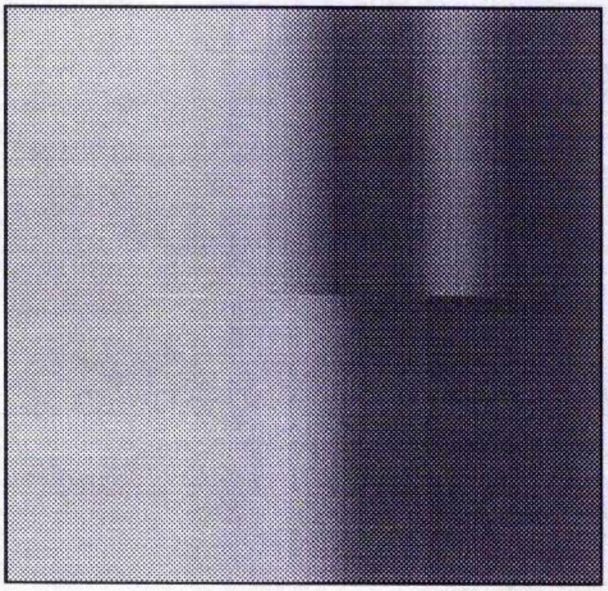

stimuli

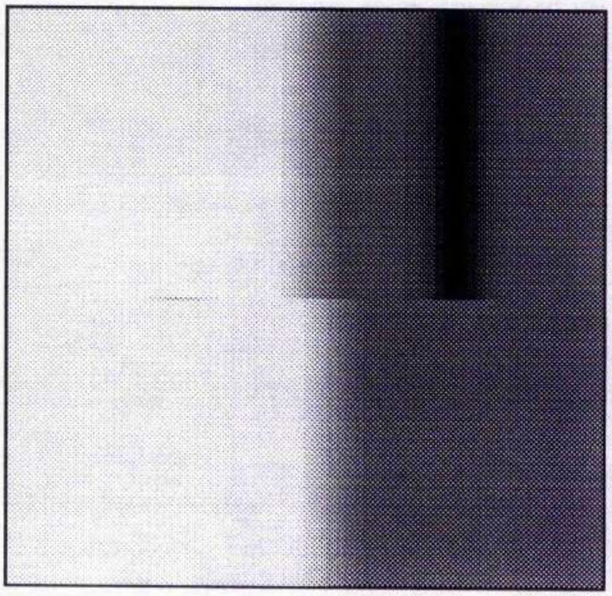

stimuli

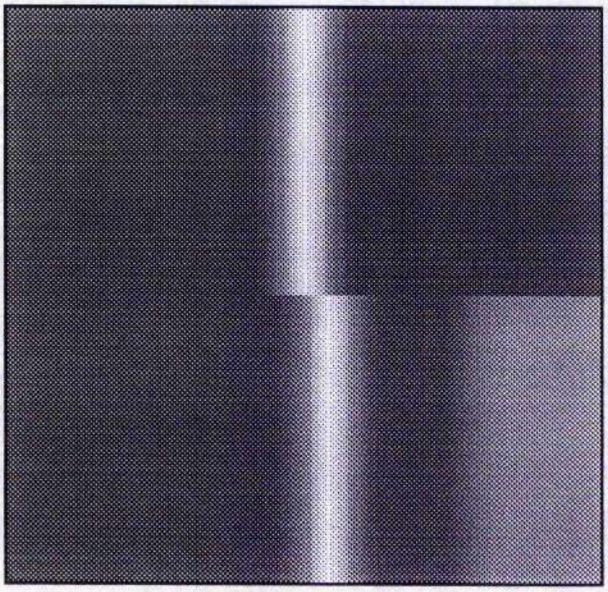

stimuli

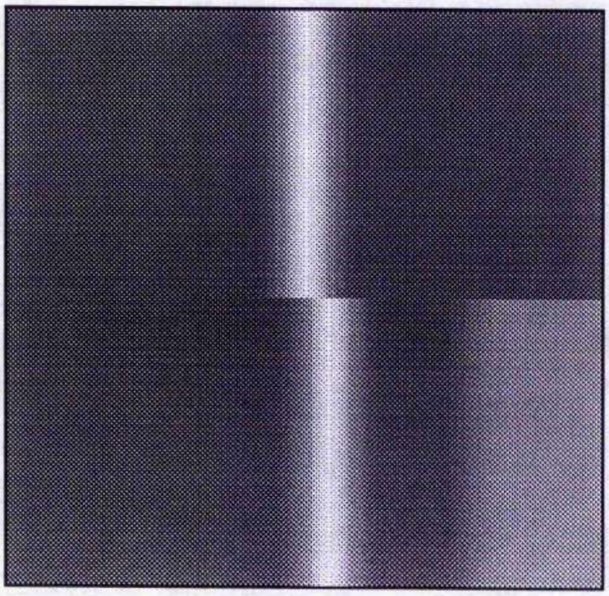

stimuli

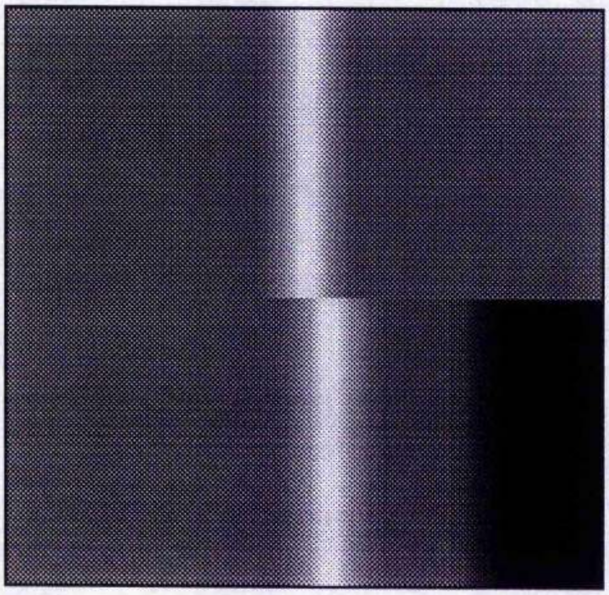

stimuli

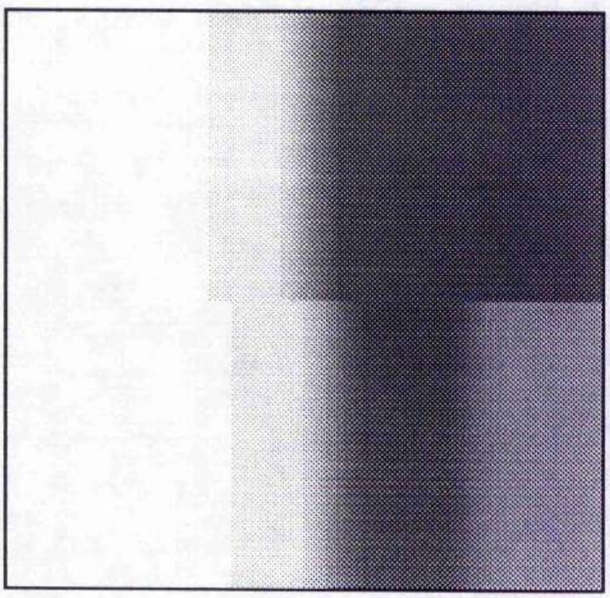

stimuli

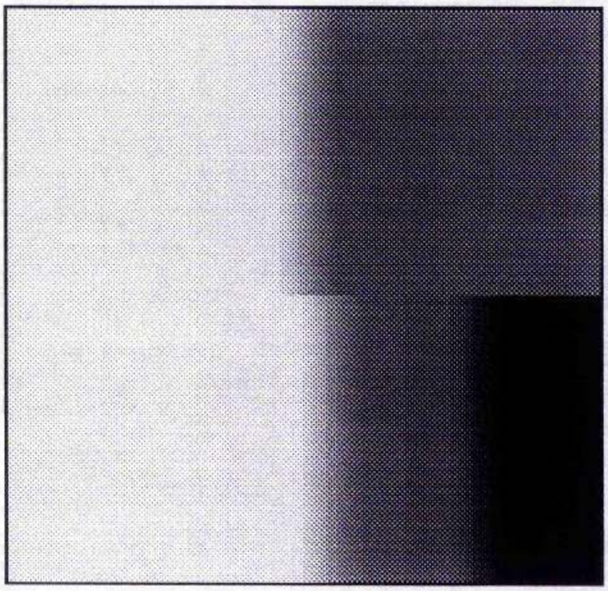

Chapter Four

Changes in perceived width of spatial intervals produced by varying the contrast profile

4.1: Overview

It is not known for certain which features determine the perceived width of a spatial interval. In this chapter, separation of two bars is defined as the separation of the luminance maxima; this is somewhat arbitrary, since any other feature in the intensity distribution (eg. the inflexions) could have been chosen. The perceived width of a spatial interval has been suggested to depend on the separation of features such as the centroids of the luminance distribution of the features delimiting the interval (Levi & Westheimer, 1987). Alternatively, the spatial separation of two luminance features may depend upon the distance between features in a neural representation of the image in which much of the information in the original intensity distribution has been lost, and information about the separation of features in the image may be encoded in the separation of spatial primitives of some description (Watt & Morgan, 1984).

The experiments reported in this section involved measuring changes in perceived width of spatial intervals caused by varying the contrast profiles of bars delimiting a spatial interval. The present experiments involved determining the point of subjective equality (PSE) of two intervals, one designated the comparison stimulus which has at least one bar with an asymmetric contrast profile; the other interval, designated the reference stimuli always had bars with a symmetric contrast profile. The first two experiments involve comparisons between stimuli which have a constant separation of the luminance inflexions of the individual bars. Systematic dependencies of the perceived width on the degree of asymmetry of the contrast profile of the bars in the comparison stimulus were obtained. The third experiment in this section measured changes in perceived width induced by varying the contrast profile of the bars delimiting the interval as in the first two, but without constraining the inflexions in the intensity profile of the bars reference and comparison stimuli to have the same fixed separation.

The results can be interpreted as evidence against features in the retinal intensity distribution being accessible to perception. Perceived separation is not equivalent to separation of the luminance maxima, threshold edges or luminance inflexions, nor can it be explained by some non-linear transform of inflexion location. The data can be accounted for by MIRAGE in parsimonious fashion, with no free parameters.

4.2: Introduction

It is still largely unknown what features in the spatial intensity distribution of the retinal image are used in the neural processing underlying spatial vision. The notion of a "local sign" or place tag (Hering, 1899; Lotze, 1884) has not been unambiguously related to the features of the neural distribution of activity consituting the visual system's reponse to a spatial pattern. It is well known that changes in perceived spatial extent can be produced by varying stimulus parameters such as contrast (Georgeson, 1980). Adaptation to a grating pattern alters perceived spatial period of gratings fairly similar in frequency (Blakemore, Nachmias & Sutton, 1970). Masking by a sinewave grating can alter the perceived size of a Difference-of-Gaussian pattern (Gelb & Wilson, 1983b). Georgeson (1980) summarised the spatial extent shifts with the observation that any manipulation which brings the stimulus nearer to detection threshold raises its apparent spatial frequency, ie. decreases its apparent spatial extent. However, there have been no generally accepted theories to explain why these size-shifts occur, and the spatial frequency channels approach has not been particularly successful; Gelb and Wilson (1983b) noted that a model based on zero-crossings in the response of a second differential filter could explain the phase-dependent size-shifts they obtained in their experiment, but a 'channels' model (assuming that size is coded by distribution of activity across spatial frequency channels) could not.

A stimulus configuration which allowed measurement of changes in the width of a spatial interval caused by manipulation of the contrast profile of bars was devised by Levi and Westheimer (1987). Levi and Westheimer (1987) varied the location of the centroid of the intensity distribution of one bar in a bar-pair interval target. They found that the perceived width of the interval followed closely the separation of the centroids of the retinal intensity distribution of the two bars. Supposing that the visual system extracts the centroid of the retinal intensity distribution for each bar cannot explain how width discrimination for a single bright bar should be possible (Watt & Morgan, 1983a). Also, the result of Levi and Westheimer for spatial interval width perception can be probably be explained by other features, such as zero-crossings in a second differential filter response (noted by Watt, Morgan & Ward, 1983b, in the context of discussing the results of Westheimer & McKee, 1977a). There is empirical evidence in favour of features in the output of second differential filters being spatial primitives in the perception of width (Watt & Morgan, 1983a). Watt and Morgan (1983b) showed that the contrast and blur dependence of vernier acuity for blurred edges and edge blur extent discrimination could be explained by supposing the visual system estimates the location of centroids of zero-bounded regions in the output of a second differential filter applied to a smoothed version of the retinal image. A study closely related to those in this chapter was conduced by Morgan, Mather, Moulden and Watt (1984), who investigated edge localisation for edges with different blurs, and explained the results in terms of a very early non-linearity in the visual system. They found that instantaneously swapping the location of edges of different blurs (whose inflexions were kept in the same location) led to apparent motion. By subjecting blurred edges to an inverse of a function used to model the intensity response function of the visual system, nulling of this apparent motion could be produced for a certain value of the parameter h in the h-transform, described below. A model involving a non-linear intensity response function preceding a second differential filter correctly explains the results obtained, by supposing that either zero-crossings or extrema (peaks and troughs) were used to represent edge location. This result indicates that any models of relative localisation for blurred features will probably have to take this non-linear transform into account. The non-linear intensity response function early in the visual system has been modelled by the following function:

$$R = \frac{(I \, . \, R_{max})}{(I + h)}$$

where $h = f(I_a)$: ie. h is a function of the adapting intensity

I = intensity,

R_{max} = half max. response of visual system,

R = response of system to stimulus of intensity I.

The h-transform has been derived from retinal physiology and fitted to psychophysical responses such as judgments about perceived brightness as a function of intensity. This transform is compressive at high intensity values. If experiments are performed where the relative location of features of different blur is compared (as in the first two experiments in this chapter), then it may be necessary to take account of the possible effects of this transform. All modelling reported in this study employed an early non-linear transform of this sort, although the parameters were fixed (using similar parameters to Toet et al, 1988, who used fairly similar stimulus conditions) rather than varied to fit the data.

On the basis of spatial interference experiments, Badcock and Westheimer (1985a) suggested that an orthoaxial region of some 3 min about a fine image feature is integrated across to obtain an estimate of the relative location of that feature.

Badcock and Westheimer demonstrated changes in the relative location of a vernier target line in the presence of an irrelevant flanking line; for small separations, the target line's location was biased towards that of the flank, and biased away from the flank at greater separations (greater than 3-4 min arc). Badcock and Westheimer explained this result by supposing that the visual system assigns location to a weighted mean of the retinal intensity distribution of a feature, with weights becoming smaller and changing sign to become negative for retinal distances greater than 3-4 min arc.

Previous experiments have explored the relation between contrast profile and relative location for bright bars. Toet, Smits, Nienhuis and Koenderink (1988) measured perceived width and perceived relative location (using vernier and interval bisection tasks) for bars with asymmetrical orthoaxial contrast profiles. In the width discrimation task, the width of an asymmetric bar was compared with that of a rectangular bar. The asymmetric bar had no internal inflexions in its luminance profile, and the distance between its outer limits was approximately 3 min arc. Perceived width of this feature was equivalent to a rectangular bar of about 2.2 min arc. Perceived width of the asymmetric bar was best explained by supposing that the visual system computes the separation of zero-crossings in the output of a second differential filter, confirming the result of Watt and Morgan (1983a). They modelled the neural response with a Laplacian of a Gaussian, which had a space constant of about 1 min arc to fit the width discrimination data. Results of the vernier task (aligning the asymmetric bar with a rectangular bar) allowed them to conclude that only zero-crossings in the filter response could correctly predict the bar's perceived relative location.

For the interval bisection task, they presented a bar with an asymmetric orthoaxial contrast profile between two rectangular bars. The bars were separated by about 10 min arc, and the outer bars were continuously visible, with the central bar being presented for 500 msec. The data was best explained by supposing that the visual system localised extrema or centroids in the response of a second differential filter. However, by explicitly instructing subjects to compare the width of the two intervals in the stimulus, they found that zero-crossings gave the best predictions of perceived relative location of the asymmetric bar. This, they suggested, indicated that the primitive employed by the visual system to solve a task was probably a function of both the spatial configuration of the target employed, and the nature of the instructions given to the subject. For width tasks, where edge-type cues may be used, they suggested zero-crossings in the filter response are extracted- for comparison of mean local signs (ie. explicit comparison of the location of luminance features), centroids or extrema in the filter response may be used. Toet et al performed their relative localisation experiments at lower contrasts, when the edges of the stimuli were not visible, but with closely similar results.

4.2.1: Aims and motivation of the experiments conducted

The first experiment in this section was designed to assess what features of the retinal stimulus controlled perceived width. Stimuli were devised for an interval width discrimination task, which consisted of two bars, one having an asymmetric orthoaxial contrast profile, and the other a symmetrical contrast profile, with the reference stimulus consisting of two symmetrical bars. Varying the degree of asymmetry of one of the bars in the comparison stimulus allowed location of features in the retinal intensity distribution such as its maxima, inflexions and centroids to be shifted. These bars were formed from sections of Gaussian luminance distribution.

The second experiment involved similar stimuli, but the focus of attention was somewhat different. The purpose of the second experiment was to determine whether there was an interaction between the spatial scale of the target, target contrast and perception of the width of the interval between two bars with asymmetric orthoaxial contrast profiles. Toet et al (1988) conducted a meticulous study of perception of relative location for bars with a skewed orthoaxial contrast profile in an interval bisection discrimination task, which is a spatial two alternative foiced choice interval discrimination. They found no consistent effect of contrast on the spatial scale of the filters which best explained the perceptual location of the bars. However, they only used fine bars, and did not vary the contrast range more than three-fold. It is possible had they used a wider contrast range or more blurred stimuli that they might have observed an interaction between the scale of filter(s) needed to explain their results and contrast. Toet et al did not examine what happens to perceived relative location of asymmetric bars if the stimuli are scaled up, and used only one degree of contrast asymmetry. In the second experiment in this section, the comparison stimulus was composed of two asymmetric bars, with the reference symmetric; two different scales of bar were used and widely separated contrasts. The third experiment here investigated the perceived width of an interval delimited by asymmetrical bars relative to an interval delimited by symmetrical bars whose space constant is the same as the inner sections (ie. between the luminance maxima of the two bars) of the asymmetric bars.

4.3: Method

4.3.1: Design

Points of subjective equality and thresholds were determined with a spatial interval discrimination task, using bars which had varying degrees of asymmetry in their orthoaxial contrast profiles. The reference stimulus always had bars with a symmetric orthoaxial contrast profile. Thresholds for stimuli with different degrees of asymmetry in the contrast profile of the bars delimiting the interval were determined in interleaved fashion.

4.3.2: Experiment 1: interval width perception, one bar delimiting interval with skewed contrast profile: stimuli

Patterns were composed of two bars. The reference stimulus consisted of two Gaussian luminance distributions. The comparison stimulus had one bar which was asymmetric and composed of two half-Gaussian distributions with different space constants, and one bar which was a Gaussian luminance distribution. The side which the asymmetric bar appeared in was varied randomly. For this experiment only, the more blurred edge was always innermost (ie. on the edge delimiting the interval). Randomly varying the side on which the asymmetric bar appeared made it less apparent to the subjects that there was anything bizarre about the comparison stimulus, which might have the effect of reducing possible biases in perception of the relative width of the stimuli.

Luminance profiles for these stimuli are given in the appendix at the end of this chapter. The sum of the space constants of the inner and outer sections of each bar was equal to twice the space constant of the reference stimulus, which had a space constant of 1.0 min arc.; eg. with an inner section blur of 1.2 min arc, the outer section had a space constant of 0.8 min arc. The inflexions in the luminance profile were thus kept a constant distance apart for all stimuli (since a Gaussian has inflexions at x=+/-s, where s is its s.d.). This meant that the separation of the inflexions in the asymmetric bar was held constant, which was a deliberate move. Watt and Morgan (1983a) had reported that the perceived width of a bright bar was largely dependent on the separation of its luminance inflexions, and keeping the separation of the luminance inflexions a constant distance apart in the present experiment was intended to ensure that the bars composing the reference and comparison stimuli were not too disparate in perceived width. Reference interval separation was either 6.0 or 9.0 min arc.

4.3.3: Experiment 2: both bars in comparison stimulus with asymmetric contrast profiles: stimuli

Patterns were composed of two bars. Each bar consisted of two halves. These halves were portions of Gaussian distributions. The two halves of the bars in the comparison stimulus were derived from Gaussian distributions which had different space constants. Luminance profiles for these stimuli are given in the appendix at the end of this chapter. The patterns presented were vertically oriented, and had a vertical extent of 45.0 min arc at 228 cm, or 180' at 57 cm viewing distance.

Note that the sum of the inner space constant and the outer space constant of each bar was always equal to twice the space constant of the reference stimulus, keeping the separation of the inflexions at a constant value for all degrees of skew of contrast profile. Unlike stimuli for the previous experiment, the edge of each bar which delimited the spatial interval could be both more and less blurred than the blur of the bars in the reference stimulus; eg. the inner blur could vary between 0.6 and 1.4 min arc for the finer stimuli, with the outer blur varying between 1.4 and 0.6 min arc. The reference stimulus had symmetric bars. Space constant of the reference stimulus was either 1.0 min arc or 4.0 min arc. A temporal 2AFC paradigm was employed; the subjects had to indicate which of two intervals presented was judged to be the wider.

Pilot trials suggested that it was possible to vary the reference interval over quite a wide range without noticeably altering the pattern of changes in perceived width obtained. It was decided to incorporate a manipulation into the experiment where the reference interval separation was varied on a pseudorandom basis from trial to trial. This forced a genuine comparison of width of the two spatial intervals, since the subject could not use just one of the stimuli in making a judgment (which would be possible if the reference interval were kept constant).

Subjects were therefore required to judge which of two intervals was the wider, with the reference interval in one trial lying within +/- 25% of the median reference interval. So on one trial, the subject might be presented with an interval of 9 and $9+\partial$ min arc, where ∂ is some change in width, and on the next trial with an interval of 14 and $14+\partial'$ min arc.

For the bars presented at 228 cm, the reference stimulus was composed of two Gaussians with space constants of 1.0 min arc. For the bars presented at 57 cm, the reference stimulus was a pair of Gaussians with space constant of 4.0 min arc. The stimuli were geometrically identical at the two viewing distances. With the reference stimulus having a space constant of 1.0 min arc, the reference interval was varied from 9-15 min arc; the reference interval on any one trial was determined by adding an amount ∂ to the median reference interval, such that $-3 < \partial < 3$ min arc (ie. 25% of the median reference interval). When the reference stimulus had a 4.0 min arc space constant, the reference interval varied over the range 36-60 min arc. The size of the increment or decrement in the spatial interval was determined by a pseudorandom number generated at runtime using a linear congruential method, which generates a uniform probability distribution.

4.3.4: Experiment 3: both bars asymmetric, inner luminance profile fixed: stimuli

Patterns were composed of two bars. Each bar consisted of two halves. These halves were portions of Gaussian distributions, truncated at the inflexions in the distribution. This truncated Gaussian was shifted downwards to ensure that the luminance returned to the background level smoothly at the point at which the distribution was truncated. The space constant of one side of this bar is defined as the distance between the luminance maximum and the zero-crossing on the appropriate side. The reference stimulus was always a symmetric bar. The two halves of the bars in the comparison stimulus were derived from Gaussian distributions which had different space constants. The inner sections of these bars delimited a spatial interval equal in luminance to the background.

The patterns presented were vertically oriented, and had a vertical extent of 40 min arc at 228 cm, 90.0 min arc at 114 cm and 180' at 57cm viewing distance. Median reference interval separation, defined as the separation between the peaks in the intensity distribution of the stimulus, was either 12.0, 24.0 or 48.0 min arc. As in experiment 2, the reference interval on any one trial was chosen pseudorandomly from the range (m +/- 25%), where m is the median reference interval. Stimuli were visible for 500 msec with an ISI of 500 msec.

4.4: Results

4.4.1: Experiment 1: Interval width perception, asymmetric/ symmetric bar pair

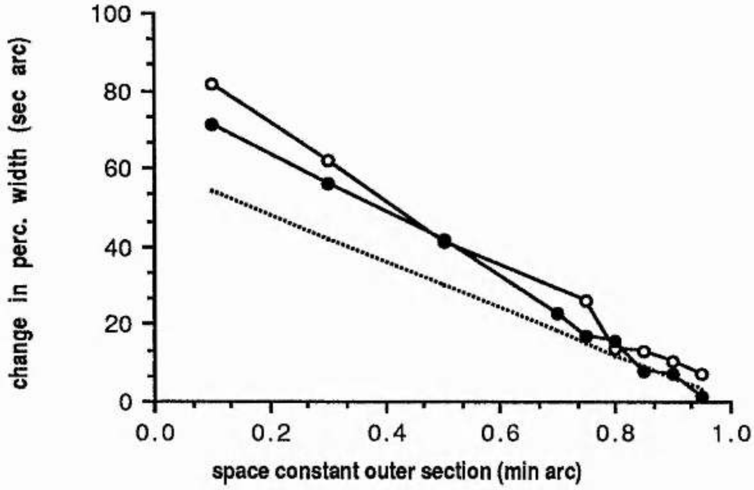

Figure 4.1: Data from a reference separation of 9.0 min arc (defined as the separation of the luminance maxima). Changes in perceived width are shown as a function of the degree of asymmetry of the contrast profile of one of the bars in the stimuli, which were pairs of bars, one of which was asymmetric, being composed of two half-Gaussian luminance distributions united at their maxima; the other bar had a Gaussian intensity profile. Changes in the perceived width of the interval between these two bars was determined relative to a reference interval delimited by two (symmetric) Gaussian bars. The Gaussian bars had space constants of 1.0 min arc. Data is shown for two observers (IRP; open circles, and GLC; filled symbols). Positive values on the ordinate indicate that the comparison stimulus was seen as narrower than the reference stimulus. The dotted line shows the location of the inflexions in the intensity profile of the comparison stimulus relative to the inflexions in the reference stimulus.

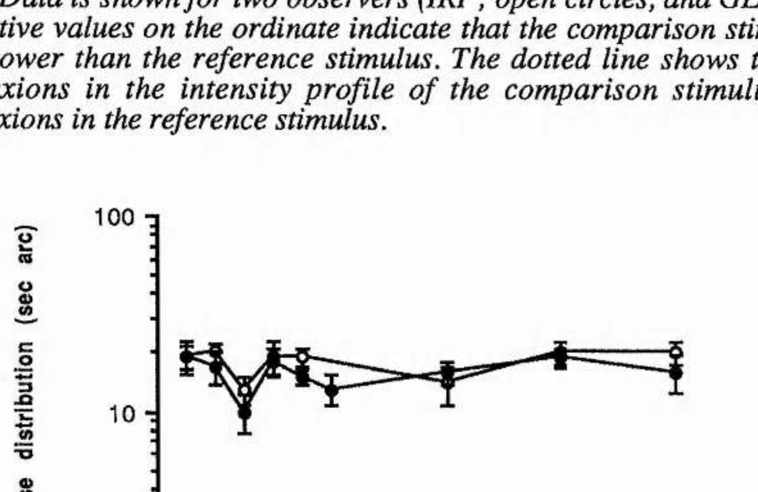

Figure 4.2:Thresholds for discrimination of interval width. Above is shown data for two subjects (IRP; open circles, and GLC; filled circles), for a reference separation (defined as the separation of the luminance maxima in the reference pair of Gaussian bars) of 9.0 min arc.

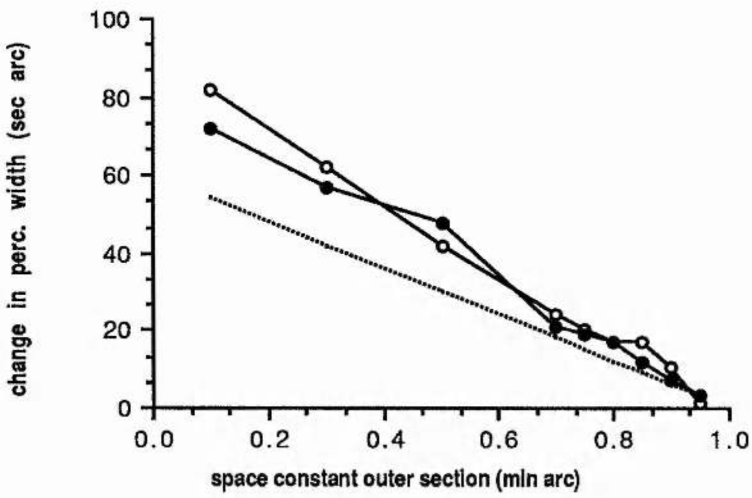

Figure 4.3: data from a reference separation of 6.0 min arc. Changes in perceived width are shown as a function of the degree of asymmetry of the contrast profile of one of the bars in the stimuli, which were pairs of bars, one of which was asymmetric, being composed of two half-Gaussian luminance distributions united at their maxima; the other bar had a Gaussian intensity profile. Changes in the perceived width of the interval between these two bars was determined relative to a reference interval delimited by two (symmetric) Gaussian bars. The Gaussian bars had space constants of 1.0 min arc. Data is shown for two observers (IRP; open circles, and GLC; filled symbols).

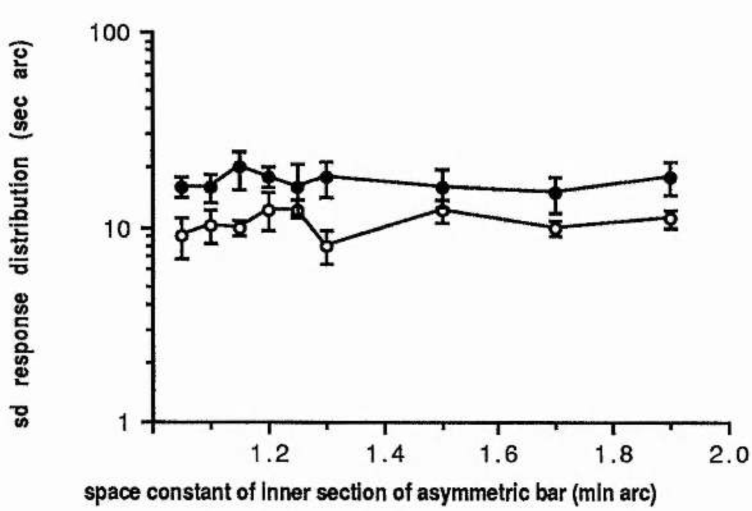

Figure 4.4: Thresholds for discrimination of interval width. Above is shown data for two subjects (IRP; open circles, and GLC; filled circles), for a reference separation (defined as the separation of tthe luminance maxima of the reference pair of Gaussian bars) of 6.0 min arc.

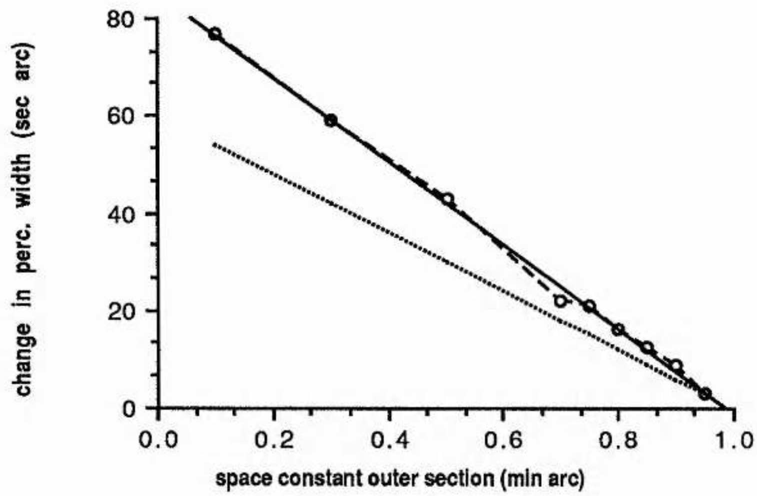

Figure 4.5: Mean shift data is shown pooled across both observers and both separations. The lower dotted line represents the shift in separation of the inflexions of the intensity profile of the comparison stimulus relative to the inflexions in the intensity profile of the reference stimulus. The continuous line is a regression line fitted to the data, which accounts for 99.8 % of the variance. The data indicate that perceived width varies more rapidly with changes in the intensity profile than changes in the locii of the luminance inflexions.

Interval thresholds and values for shifts in the perceived width of the spatial interval delimited by the two bars are presented above, as a function of the space constant of the inner section of the one bar in the target which is asymmetric. There is good agreement between the subjects, and the mean shift data is similar at both reference separations. Data pooled across the subjects and reference separations is also shown above.

Separation of the bars is defined as the separation of the luminance maxima. The data show that perceived separation is not equivalent to separation of luminance maxima. The positive values on the ordinate indicate a decrease in perceived width of the comparison interval relative to the reference interval width (equivalent, in a sense, to perceptual attraction, as reported in experiment 6 in chapter 3 for interval width perception, and by Levi & Westheimer, 1987).

The data indicate that as the interior of the asymmetric bar becomes more blurred, there is a gradual decrease in the relative perceived separation of the interval delimited by this asymmetric bar and another symmetric bar. Perceived appearance of the individual bars was kept relatively constant, by maintaining a constant distance between the luminance inflexions.

4.4.2: Experiment 2: Interval width perception, asymmetric bar pair

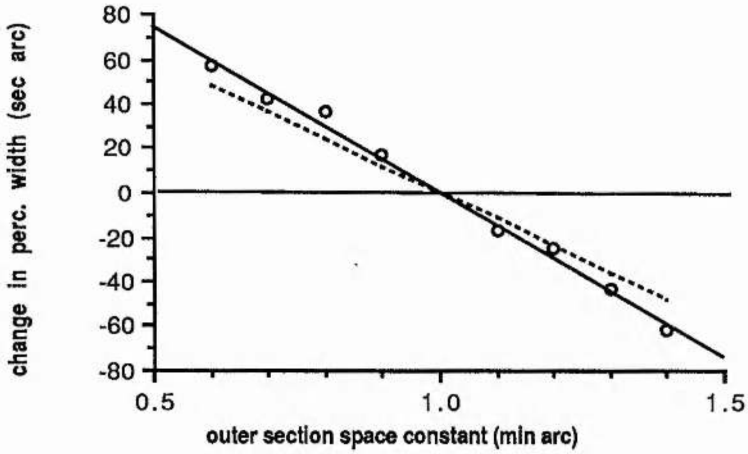

Figure 4.6 : Changes in the perceived separation of two bars, caused by varying the degree of asymmetry of the orthoaxial contrast profile; this was collected at the maximum contrast available. The data shown has been pooled across the two observers (GLC, IRP). The reference stimulus for this task was a pair of Gaussian bars with space constants of 1.0 min arc. Median reference separation of the luminance maxima was 12 min arc. A negative value on the ordinate indicates that the interval delimited by the (asymmetric) comparison stimulus was seen as narrower than the reference interval. The dotted line represents the change in location of the inner inflexions in the intensity profile of the asymmetric stimulus, relative to the inflexions of the reference stimulus. The continuous line is a regression line fitted to the data.

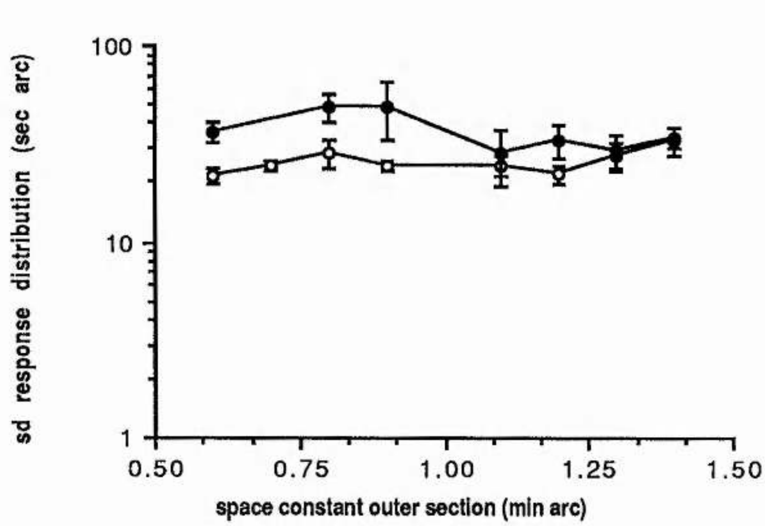

Figure 4.7: Interval thresholds are shown for two subjects (IRP; open circles, and GLC; closed circles). Median separation was 12 min arc jittered with a 6 min arc range. and the reference stimulus for this task was a pair of Gaussian bars with space constants of 1.0 min arc.

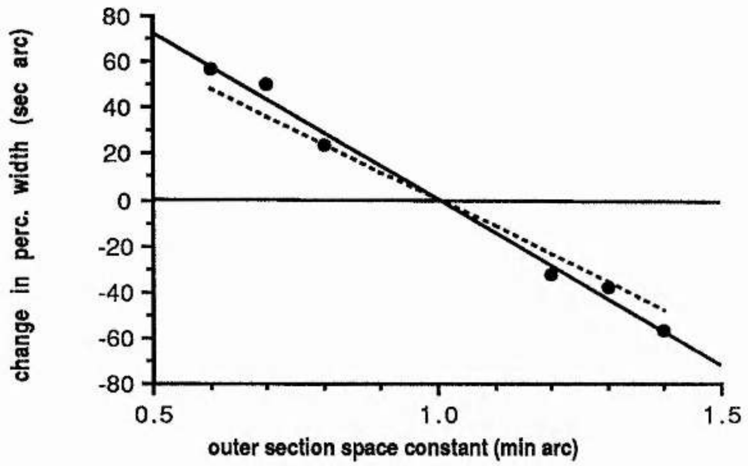

Figure 4.8: Data from a contrast 15dB below the maximum contrast used above (ie. Michelson contrast of 0.08). The stimuli were otherwise identical, and the data has been pooled across the two observers (GLC, IRP). The reference stimulus for this task was a pair of Gaussian bars with space constants of 1.0 min arc. Median reference separation was 12 min arc, jittered with a 6 min arc range. The dotted line is the relative location of the inflexions in the comparison stimulus, and the continuous line a regression line fitted to the data.

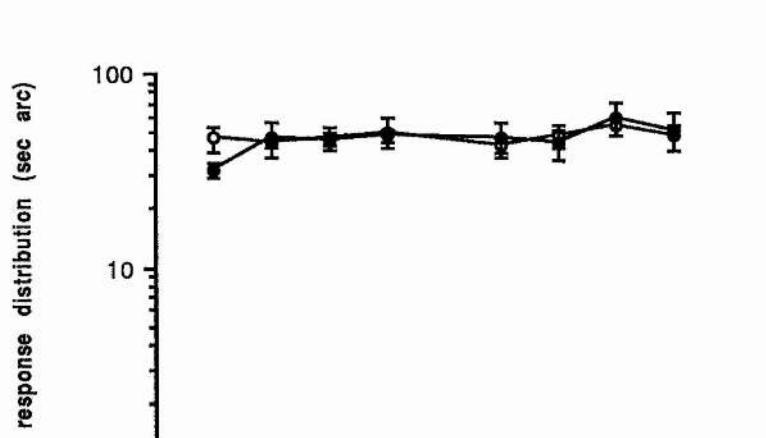

Figure 4.9: Interval thresholds are shown for two subjects (IRP; open circles, and GLC; closed circles) for the lower contrast (15dB below the maximum available) data above. Median separation was 12 min arc, jittered with a 6 min arc range and the reference stimulus for this task was a pair of Gaussian bars with space constants of 1.0 min arc.

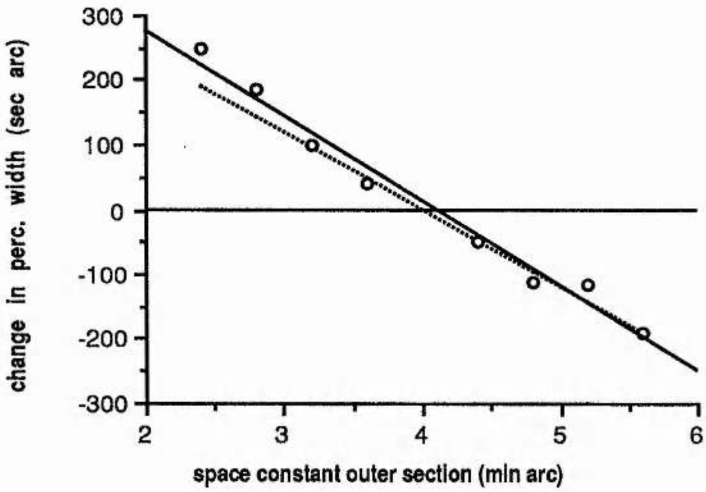

Figure 4.10: Changes in the perceived separation of two bars, caused by varying the degree of asmmetry of the orthoaxial contrast profile; this was collected at the maximum contrast available. The data shown has been pooled across the two observers (GLC, IRP). The reference stimulus for this task was a pair of Gaussian bars with space constants of 4.0 min arc, with a median reference separation of 48 min arc, jittered over a 24 min arc range. The dotted line represents the change in location of the inner inflexions in the intensity profile of the asymmetric stimulus, relative to the inflexions of the reference stimulus. The continuous line is a regression line fitted to the data.

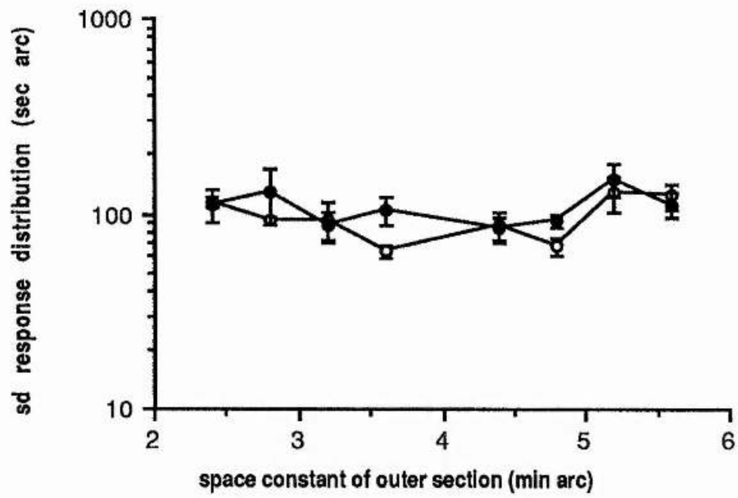

Figure 4.11: Interval thresholds are shown for two subjects (IRP; open circles, and GLC; closed circles). Median separation was 48 min arc, and the reference stimulus for this task was a pair of Gaussian bars with space constants of 4.0 min arc, presented at the maximum contrast available.

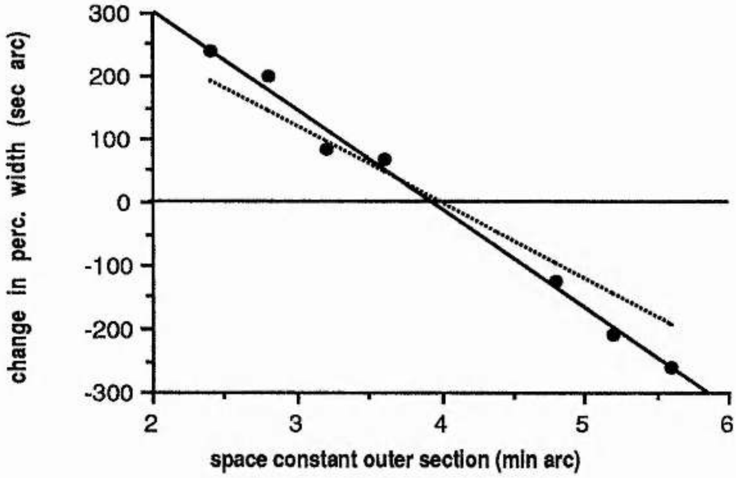

Figure 4.12: Changes in the perceived separation of two bars, caused by varying the degree of asmmetry of the orthoaxial contrast profile, collected at 25dB below the maximum contrast available (ie. 0.025). The data shown has been pooled across the two observers (GLC, IRP). The reference stimulus for this task was a pair of Gaussian bars with space constants of 4.0 min arc, with a median reference separation of 48 min arc. The dotted line represents the change in location of the inner inflexions in the intensity profile of the asymmetric stimulus, relative to the inflexions of the reference stimulus. The continuous line is a regression line fitted to the data.

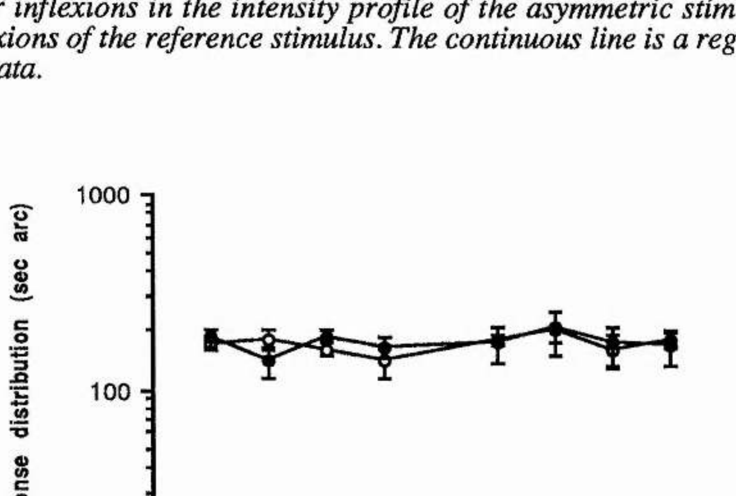

Figure 4.13: Interval thresholds are shown for two subjects (IRP; open circles, and GLC; closed circles), for a contrast 25dB below the maximum available. Median separation was 48 min arc, and the reference stimulus for this task was a pair of Gaussian bars with space constants of 4.0 min arc. Precision is not affected by the degree of asymmetry of the intensity profile.

Interval thresholds and changes in the width of an interval between two asymmetric bars, relative to the width of an interval between two symmetric bars are presented above in figures 4.6-4.13 as a function of the space constant of the outer section of the comparison stimulus. The data has been pooled across both subjects, since their data was very similar. Note that a negative value on the ordinate indicates that the comparison stimulus was seen as wider than the reference stimulus.

The stimuli used in this experiment are very similar to those used above, with the difference that there are two rather than just one asymmetric bar(s). Separation of the luminance inflexions of the bars is kept constant, to maintain the bars at approximately equal perceived widths; the bars with 1.0' Gaussian bars as a reference all have a 2.0' separation between their luminance inflexions, and the bars with a 4.0' reference have an 8.0' inflexion separation. Although two different contrasts were tested at, the perceived width of intervals was not compared across different contrasts within a session, ie. the data was collected at constant contrast.

The first experiment looked only at stimuli where the inner section of one bar delimiting the interval increased in blur. In this experiment, the outer section of both bars were varied in blur, so that the outer section could be either more blurred or less blurred than the inner section. The perceived width of a spatial interval delimited by two bars with asymmetrical orthoaxial contrast profiles changes relative to an interval delimited by two symmetric bars varies as the degree of asymmetry of the contrast profile is altered. As the blur of the outer section increases, the perceived width of the interval increases.

It can be seen that the relation between the skew of the bars delimiting the interval and the relative perceived width is not obviously dependent on contrast. The differences in contrast introduced were quite large, but the changes in perceived position as a function of the degree of asymmetry of the contrast profile were relatively similar at both high and low contrasts.

4.4.3: Experiment 3: inner section width constant, outer section space constant varied

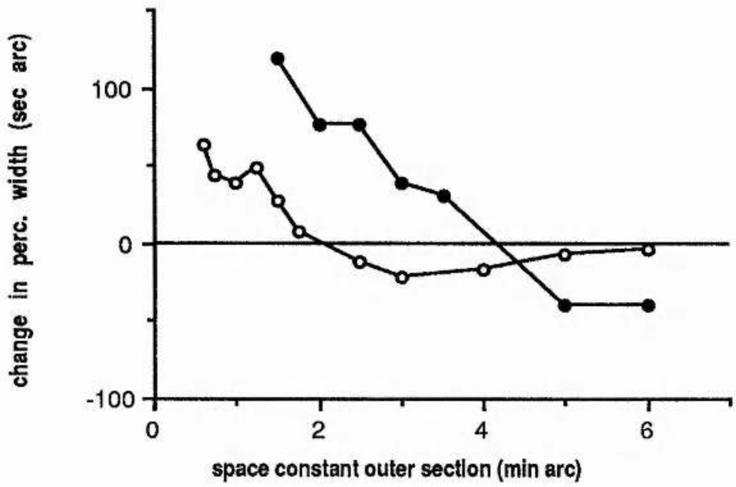

Figure 4.14: Data from two subjects (IRP, GLC) has been combined, since they displayed closely similar results. Open circles are from the condition where the inner section of the two bars had a space constant of 2.0 min arc (with a median reference separation of 12.0 min arc jittered over a range of 6 min arc), and the filled circles are from the condition where the inner section of each bar had a space constant of 4.0 min arc (with a mean reference separation of 24.0 min arc and a range of 12 min arc). A positive value on the ordinate indicates that the comparison interval (with both bars asymmetric) is seen as narrower than the reference interval, which is delimited by symmetric bars.

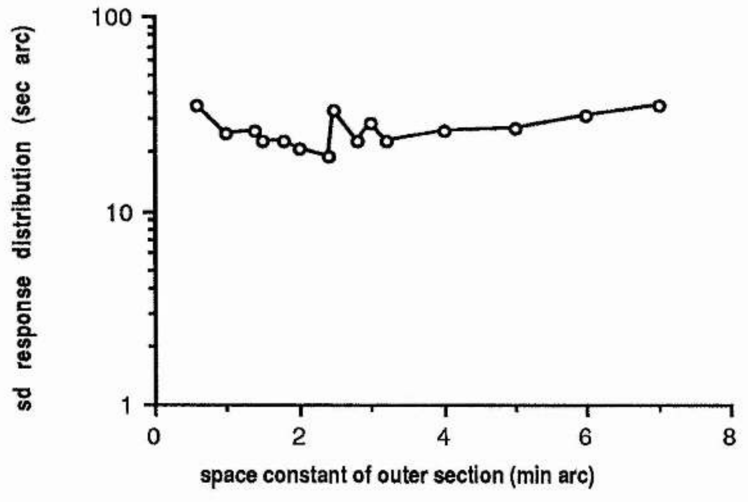

Figure 4.15: Thresholds are shown for subject IRP, for the condition where the inner section had a space constant of 2.0 min arc, and the median separation was 12 min arc.

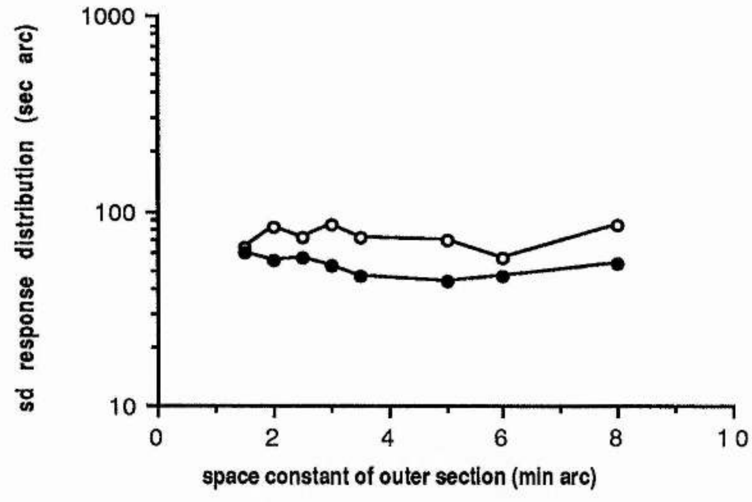

Figure 4.16: Thresholds are shown for both subjects (IRP; open circles, and GLC; filled circles), for the condition where the inner section had a space constant of 4.0 min arc, and the median separation was 24 min arc.

Interval thresholds and changes in perceived width of the interval between two asymmetric bars relative to the width of an interval delimited by two symmetric bars are presented above in figures 4.14-4.16. Data shown is the mean of two subjects (IRP and GLC), whose results were closely similar. Whilst the inner section of each bar remains constant throughout, the width of the outer section of the bar has some influence on the perceived width of the interval at least within the parameter ranges investigated here. The function relating the changes in perceived width to outer section space constant is non-monotonic, indicating that beyond a certain width of outer section, increasing the width leads to a decrease in the perceived width of the interval between the bars.

Note that the inner section profile of the reference and comparison stimuli is identical throughout. Hence the inner inflexions, luminance maxima and up to 4 min arc (for the wider stimuli) of the luminance profile adjacent to the spatial interval does not change between the reference and comparison stimuli, but the reference and comparison intervals are perceived as being of different spatial extents, and this depends systematically on the degree of asymmetry in the contrast profile of the bars delimiting the interval.

4.5: Discussion

4.5.1: Interval width perception and the retinal intensity distribution

Whilst the peak of the stimulus intensity distribution remains in the same place throughout the experiments reported here, the perceived separation of the target bars changes systematically. The peak of the stimulus intensity distribution is apparently not available to perception, which has been observed by previous authors (Watt & Morgan, 1983a). Inflexions in the stimulus intensity distribution are not the physical delimiters of perceptual intervals. This is apparent from the data from experiment 1, where the perceived separation is much less than would be expected if the luminance inflexions were perceptual delimiters of the spatial interval. The comparison stimulus for that experiment had one bar with an edge (adjacent to the interval) more blurred than either of the two edges in the reference stimulus; an inflexions model would understimate the change in interval width induced by manipulating the contrast profile of one bar. One explanation of this could be some early non-linearity in transduction of intensity information, which has been used by Morgan, Mather, Moulden and Watt (1984) to explain why apparent motion is seen when edges of different blurs are aligned by their luminance inflexions. On the other hand, this explanation is not compatible with the data from experiment 3, where the inner profile of the bars remain constant throughout, but the perceived width changes systematically. An inflexions model with a non-linearity cannot explain why the perceived width should change if the inner section of each bar delimiting the spatial interval does not change. The data from the third experiment showed that even when the whole of the inner section of the bars delimiting the interval is held constant, the perceived width of the interval depends on the intensity profile outwith the luminance maxima; ie. the inner inflexions in the retinal intensity profile are not perceptual delimiters of spatial intervals. The dependence of perceived width of a spatial interval on luminance information beyond the luminance maxima (ie. the outer section of the bars) indicates that extraction of some intensity feature of the inner section of the bars delimiting the interval (eg. threshold edges) cannot explain the results (confirming results of Watt & Morgan, 1983a and Toet, Smits, Nienhuis & Koenderink, 1988).

4.5.2: Spatial primitives and interval width perception

In the first experiment, where the inner edge of each bar was blurred, the perceived width of the interval delimited by the asymmetric bar and its symmetric counterpart was narrower than predicted by the separation of the inflexions in the retinal intensity distribution. This is unlikely to be due to an early non-linearity, given the results of the third experiment. The data is well-accounted for by a model involving localisation of centroids in the output of a MIRAGE algorithm (with filters having space constants from 0.35-2.8 min arc), or by the zero-crossings in a filter with a space constant of around 1.0 min arc.

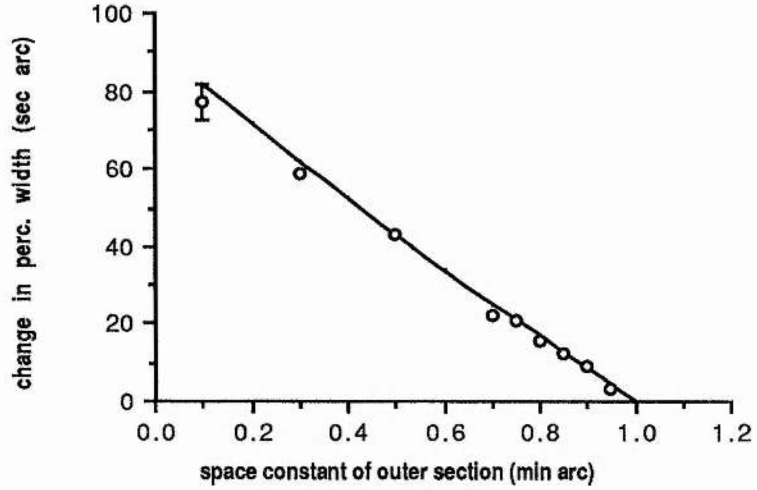

Figure 4.17: Data and predictions from the first experiment, involving measuring location biases for stimuli with one of two bars having an asymmetric orthoaxial contrast profile. The above graph shows predictions of MIRAGE, assuming that the bars are represented at the location of the centroids in the positive portion of the output of the MIRAGE algorithm for filters with space constants from 0.35-2.8 min arc. The open circles represent the pooled data (pooled across both observers and both separations).

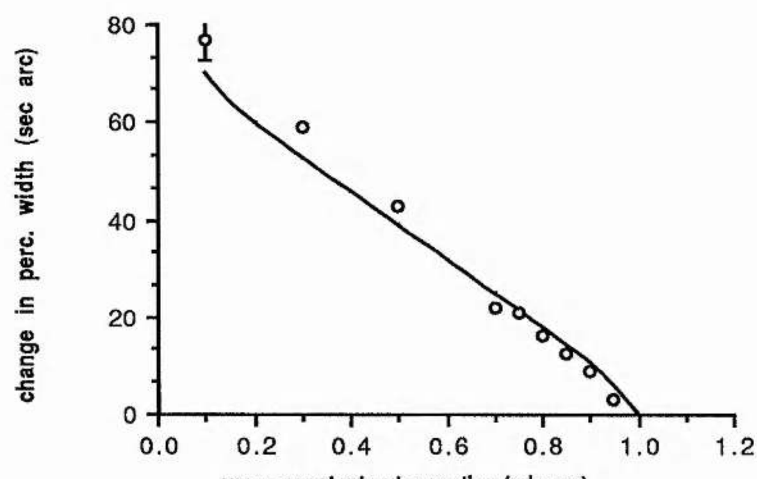

Figure 4.18: Data and further predictions from the first experiment,. The above graph shows predictions of a model involving zero-crossings in a single scale of filter (a LoG filter with a space constant of 1.0 min arc), assuming that the separation of the bars is judged as the separation of the inner zero-crossings in the filter response. The open circles represent the pooled data (pooled across both observers and both separations).

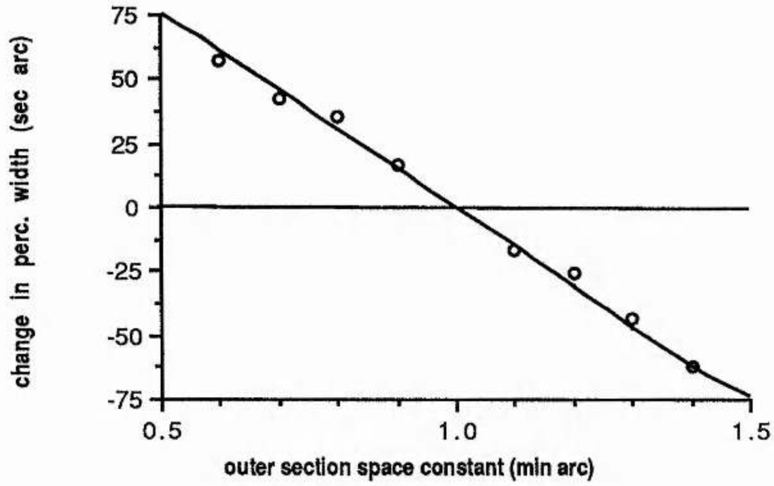

Figure 4.19: Experiment 2; data from discrimination of separation of asymmetric bar pairs, at maximum contrast. Data (open circles) and predictions of MIRAGE (continuous line), when the centroids in the positive part of the filter response are used to represent the location of the bars. The space constants of the filters used (LoG's) ranged from 0.35-2.8 min arc, after Watt and Morgan (1984).

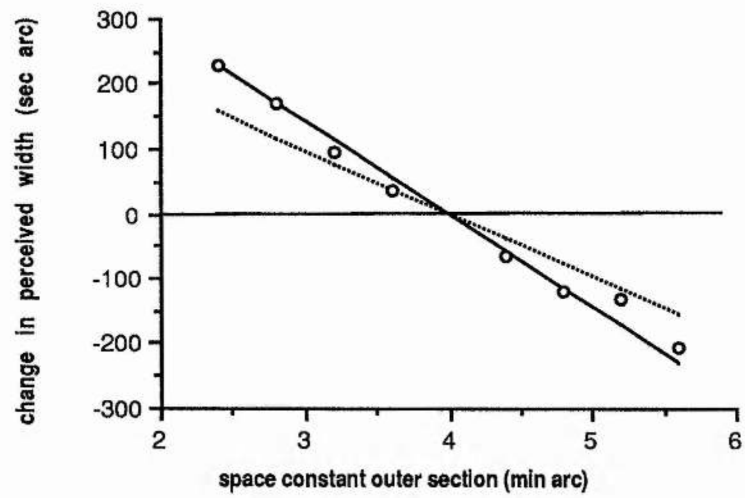

Figure 4.20: Experiment 2; data from discrimination of separation of asymmetric bar pairs, at maximum contrast. Reference stimulus for this task was a pair of Gaussian bars with a space constant of 4.0 min arc. Data (open circles) and predictions of MIRAGE, when the centroids in the positive part of the filter response are used to represent the location of the bars. The space constants of the filters used (LoG's) ranged from 0.35-5.6min arc for the continuous line, and 0.35-2.8 min arc for the dotted line. The data is best fitted by supposing recruitment of a larger filter than 2.8 min arc.

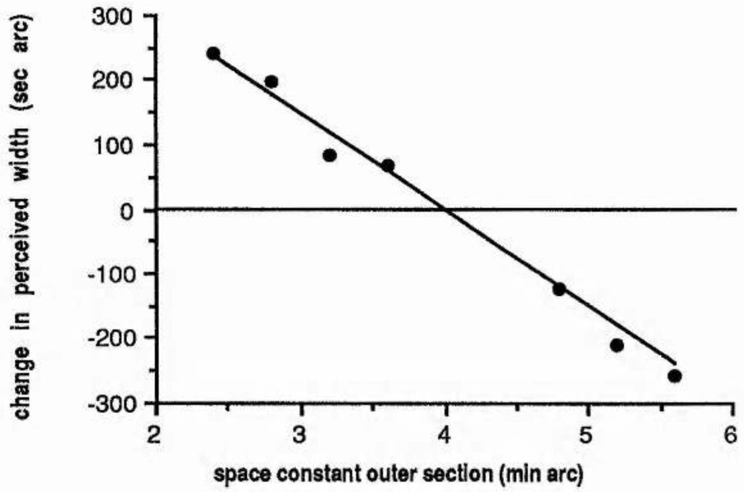

Figure 4.21: Experiment 2; data from discrimination of separation of asymmetric bar pairs, at 25dB below contrast. Reference stimulus for this task was a pair of Gaussian bars with a space constant of 4.0 min arc. Data (filled circles) and predictions of MIRAGE, when the centroids in the positive part of the filter response are used to represent the location of the bars. The space constants of the filters used (LoG's) ranged from 0.35-5.6min arc. The data is well fitted by this model, implying a similar relative localisation strategy is used across an 18-fold range in contrast (25dB=1.25 log units).

One explanation of the data in the first two experiments of this chapter is to suppose that perceptual separation is controlled by the location of centroids in the output of something akin to a MIRAGE operation. Note that MIRAGE was successful in explaining the data from spatial interference experiments in chapter three if extrema were localised; however, if stimuli are very close together, the output of MIRAGE does not yield two response distributions which are distinct from one another, and only the peaks in the response distribution can be used to assess the location of closely spaced bars (ie. centroids as primitives require distinct regions of response of the same sign to the two bars). It is therefore reasonable that a switch in localisation strategy occurs if two distinct regions of response can be generated in response to two nearby bars, with a change from localising extrema to localising centroids. There are known differences between separation discrimination for separations above and below about 5 min arc; Levi and Westheimer (1987) found that below this separation, thresholds were increased if the bars delimiting the interval were of opposite polarity of contrast (ie. bright/ dark is worse than bright/bright or dark/dark), but beyond this separation, there is no effect of the polarity of contrast on interval discrimination thresholds. Additionally, Morgan and Regan (1987) found that contrast randomisation of line targets did not affect spatial interval discrimination unless the interval was as narrow as 2.5 min arc.

Conceivably, the retinal distance over which spatial interference is greatest (4-5 min arc; Westheimer & Hauske, 1975) may reflect a transition between these two strategies, with neither being employed optimally- by analogy with mesopic vision, when neither the rod nor cone systems function at normal levels of efficiency, and many aspects of visual performance are poor.

The data from the first two experiments here is well explained by supposing that the visual system is localising centroids int he output of a MIRAGE operation, although it is not conclusive, since it can equally well be explained by looking at the zero-crossings in an appropriately chosen scale of spatial filter. To explain the data from the coarser scale of bars used in the second experiment, an additional LoG filter with a space constant of 5.6 min arc is necessary; without this additional filter, the predicted effects of the contrast profile manipulation are too small. The data from the different contrasts in the second experiment was quite similar, in terms of the pattern of changes in perceived interval width obtained, and can be explained by MIRAGE without any necessity to change the range of filters involved.

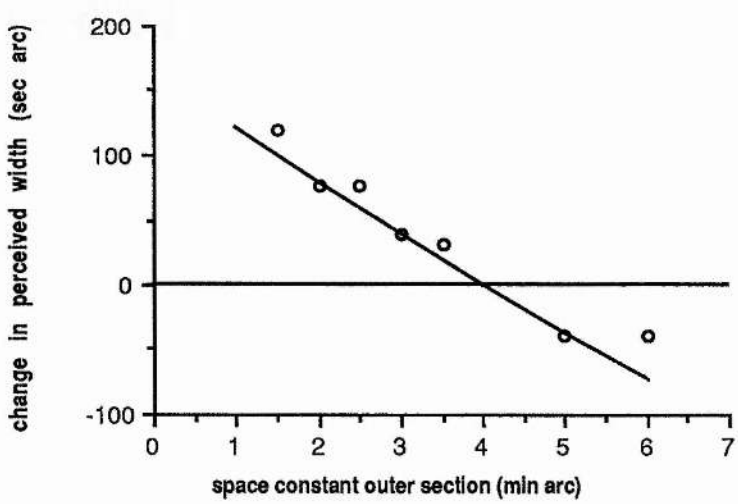

Figure 4.22: Modelling of data from the third experiment, where the reference stimulus consisted of symmetric, truncated Gaussian bars, and the comparison stimulus was asymmetric. The inner section space constant was fixed throughout. Data from the condition where the separation was 24.0 min arc and the inner section space constant 4.0 min arc, plus predictions of MIRAGE., are shown above. The results of applying the same model as previously, involving determining the separation of centroids in the output of a MIRAGE algorithm, with LoG filters having space constants from 0.35-2.8 min arc are illustrated. The data is apparently quite well fitted by this model.

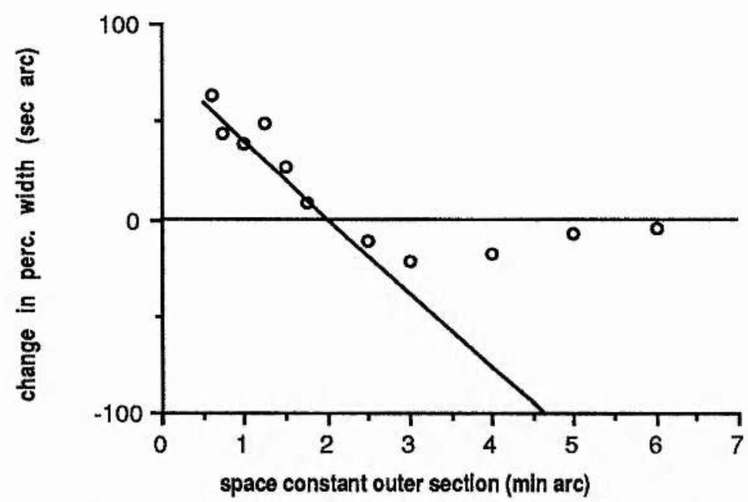

Figure 4.23: An indication of MIRAGE breaking down. Data from the condition where the separation was 12.0 min arc and the inner section space constant 2.0 min arc, plus predictions of MIRAGE. Results of applying the same model as above, involving determining the separation of centroids in the output of a MIRAGE algorithm, with LoG filters having space constants from 0.35-2.8 min arc, to the data from the third experiment; MIRAGE would predict a monotonic increase in perceived width as the space constant of the outer section is increased, but this is not the case. However, the data is quite well fitted by this sort of model for the smaller values of space constant for the outer section.

In the third experiment, it was observed that perceived width could be varied by altering the space constant of the outer section of the bars, although perceived width was not a monotonic function of the width of the outer section, and began declining if perceived width became too large. There are some difficulties in the interpretation of this experiment in terms of simple ideas about spatial primitives, since the reference and comparison stimuli looked quite different. This difference in appearance may have introduced other biases in perception; for instance, the context dependency of spatial extent perception is well known; the Baldwin illusion, for example, refers to an apparent decrease in spatial interval between two bars if the width of the bars is increased. Whatever processing generates this illusion as a side-effect, it probably lies at a stage of visual processing beyond extraction of spatial primitives. Therefore, it seems inappropriate to attempt to explain this in terms of spatial primitives. However, the change in perceived width observed in this third experiment using similar stimuli to the first two experiments in chapter four indicates that assumptions made about changes in perceived width were correct; the perceived width depends on a moderately large region of the intensity profile of the stimulus (ie. greater than 4 min arc even for these truncated bars with quite sharp edges) and changing the degree of asymmetry of the intensity profile of bars can alter the perceived width of intervals delimited by these bars.

4.6: Summary

(1) Changes in perceived width, without accompanying changes in the precision of interval discrimination, were measured using spatial intervals delimited by two bright bars composed of sections of Gaussian luminance distributions with different degrees of asymmetry in the orthoaxial contrast profile. Skewing the contrast profile outwards (ie. away from the interval between the bars) increases the perceived width of the interval, and skewing it inwards decreases the perceived width. The changes in perceived width cannot be explained by features in the intensity distribution (luminance maxima, inflexions, threshold edges) and an early non-linear transform of the location of inflexions in the retinal intensity distribution cannot explain the observed changes in perceived width, since the perceived width can be altered by varying the skew of the contrast profile even when the inner sections of the bars delimiting the spatial interval remain constant. Varying the degree of asymmetry has a similar effect across a wide range of contrasts (up to 18-fold for the coarsest scale of stimuli used).

2) MIRAGE accounts for the data by supposing centroids output of MIRAGE are localised. The predictions of MIRAGE break down if the reference and comparison stimuli are too different in appearance (experiment 3), suggesting that higher-level processes must be invoked in this situation.

Appendix: Stimuli from chapter 4

Experiment 1: one bar in comparison stimulus with asymmetric contrast profile

Luminance profiles for the stimuli were given by:

$$L(x) = L_b[1 + c.w(x)]$$

where

L_b = background luminance,
c = contrast (0.45),

w(x) is identical to that given below for the stimuli in experiment 2, with the slight difference that only one of the two bars in the comparison stimulus was asymmetric. Examples of this function can be seen overleaf, as well as greyscale representations of some of the images used. Recommended viewing distance for the greyscale images is **5 metres**.

Experiment 2: both bars in comparison stimulus with asymmetric contrast profiles

Luminance profiles for the stimuli are given by:

$$L(x) = L_b[1 + c.w(x)]$$

L_b = background luminance,
c = contrast,

For the comparison stimulus,

$$w(x) = A\ (\ \exp[\frac{-(x-\mu_1)^2}{2s_1^2}] + \exp[\frac{-(x-\mu_2)^2}{2s_2^2}]\ )$$

where μ_i = mean position of i^{th} bar,
s_i = space constant of original Gaussian distribution,
A is a scaling constant,

and given $\mu_2 > \mu_1$, for $(x-\mu_1) < 0$, s_1 = outer width of 1st bar,
for $(x-\mu_1) > 0$, s_1= inner width of 1st bar
and likewise for $(x-\mu_2) < 0$, s_2= inner width of second bar
for $(x-\mu_2) > 0$, s_2 = outer width of 2nd bar,
$\mu_2 - \mu_1$ = peak-to-peak separation of luminance maxima of the bars

Examples of this function and greyscale representations of the images used can be seen overleaf.

Experiment 3: both bars asymmetric, inner section space constant held fixed

Luminance profiles for the truncated Gaussian stimuli are given by:

$L(x) = L_b[1 + c.w(x)]$
L_b = background luminance,
c = contrast,

For the comparison stimulus,
$w(x) = A\ [\ g(x) + h(x)\ ],$
where

$$g(x) = \exp[\frac{-(x-\mu_1)^2}{2s_1^2}] - \exp(-0.5)$$

$$h(x) = \exp[\frac{-(x-\mu_2)^2}{2s_2^2}] - \exp(-0.5)$$

with g(x) and h(x) strictly greater than 0
where μ_i = mean position of i^{th} bar,
s_i = space constant of original Gaussian distribution,
A is a scaling constant,
and given $\mu_2 > \mu_1$, for $(x-\mu_1) < 0$, s_1 = outer width of 1st bar,
for $(x-\mu_1) > 0$, s1= inner width of 1st bar (1.0 -2.0 min arc)
and likewise for $(x-\mu_2) < 0$, s_2= inner width of second bar (1.0-2.0 min arc)
for $(x-\mu_2) > 0$, s2 = outer width of 2nd bar
$\mu_2 - \mu_1$ = peak-to-peak separation of luminance maxima of the bars

w(x) for the reference stimulus is identical in form, with the exception that s_1=s_2 and the bars are symmetrical (ie. s_i does not depend on x).

The constant exp[-0.5] appears in the equations for g(x) and h(x) above, since this represents the value on the ordinate for a Gaussian distribution at its inflexions, ie. at one standard deviation above and below the mean. Its inflexions are located at the zeros in its second derivative, which is given by the formula

$$f''(x) = k\ (\ \frac{x^2}{s^2} - 1)\ f(x)\ ,$$

where f(x) is a Gaussian function. This has zeroes at x = s, x = -s .

Several different examples of the above function can be seen overleaf, as well as greyscale representations of the images used. Also included is a version in which the space constant of the outer section varies from top to bottom of the greyscale image. This demonstrates a new illusion, in which the shape of the interval varies as a function of viewing distance. From the correct viewing distance (ie. > 5 metres), the interval appears to be bowed outwards, ie the width increases non-monotonically as one moves down the image. At smaller viewing distances, there is an unambiguous increase in width as one moves down the image.

Chapter 4 stimuli

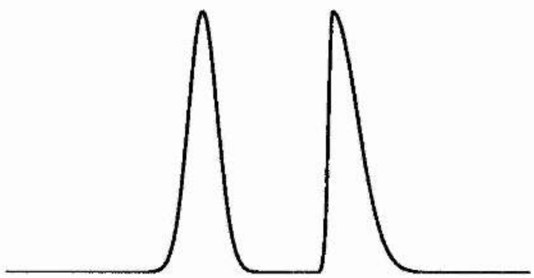

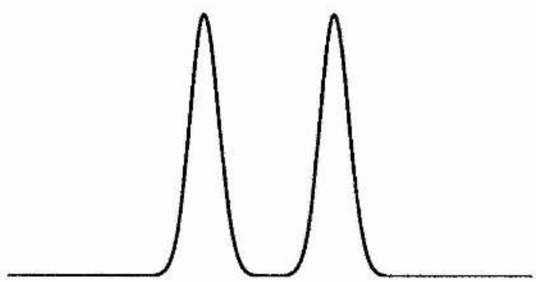

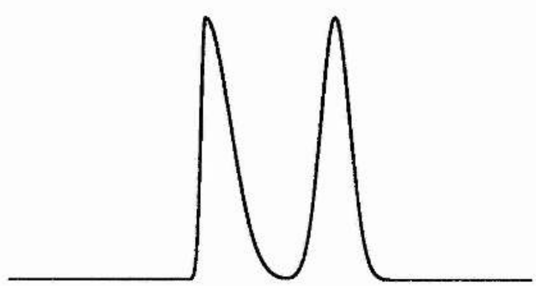

Chapter 4 stimuli

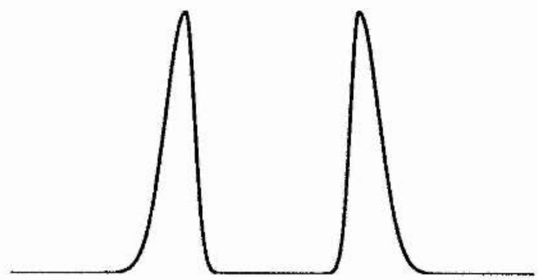

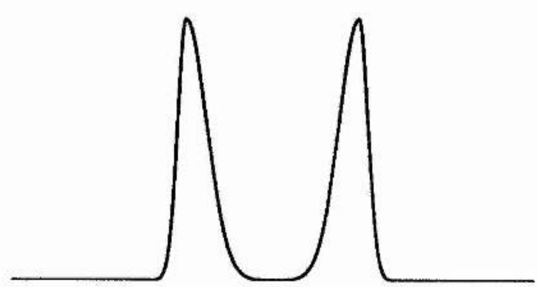

stimuli

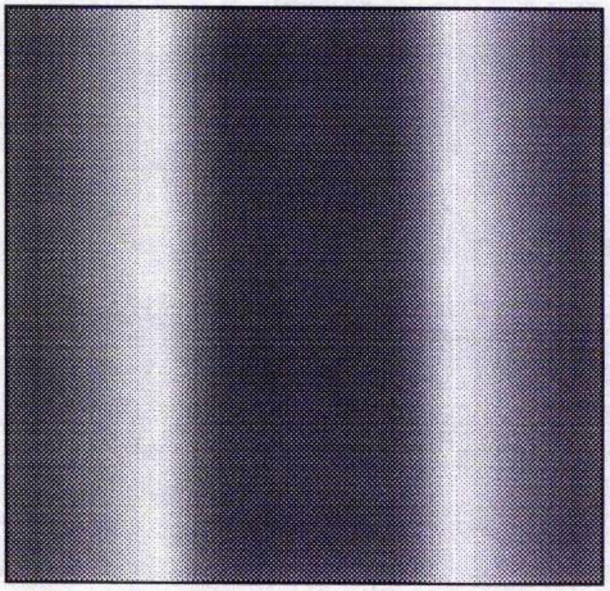

stimuli

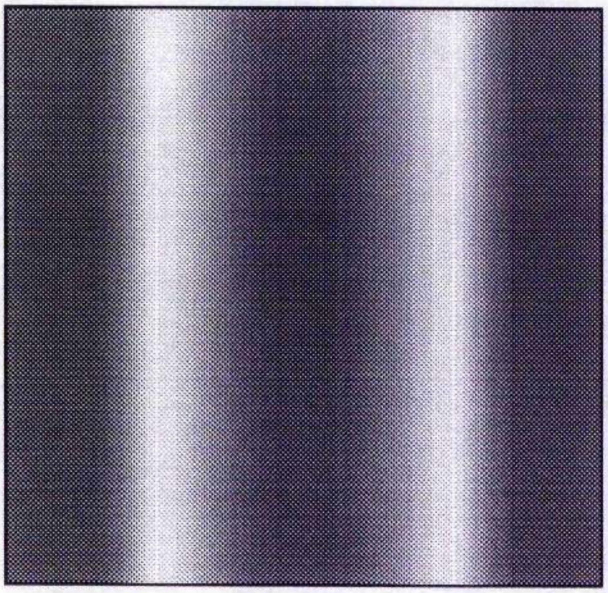

stimuli

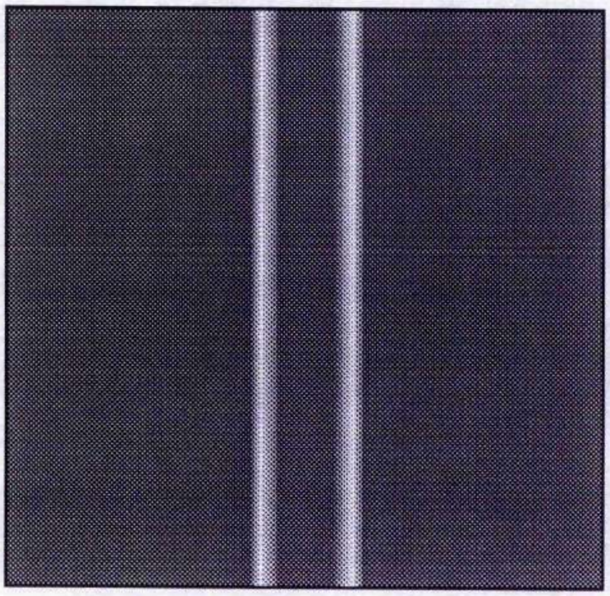

stimuli

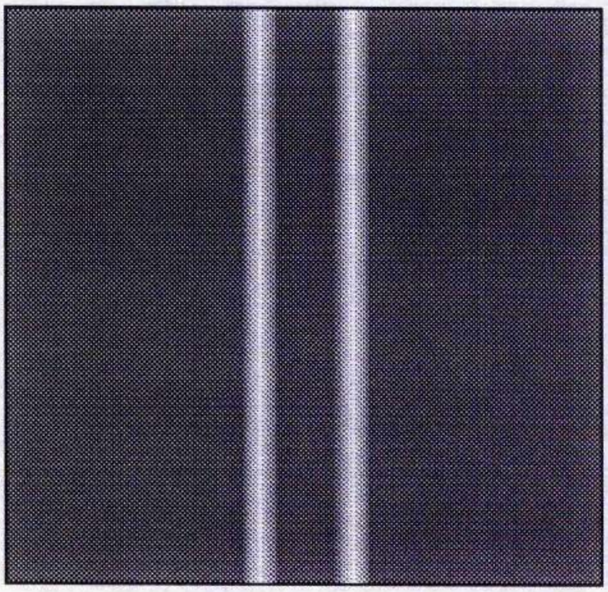

stimuli

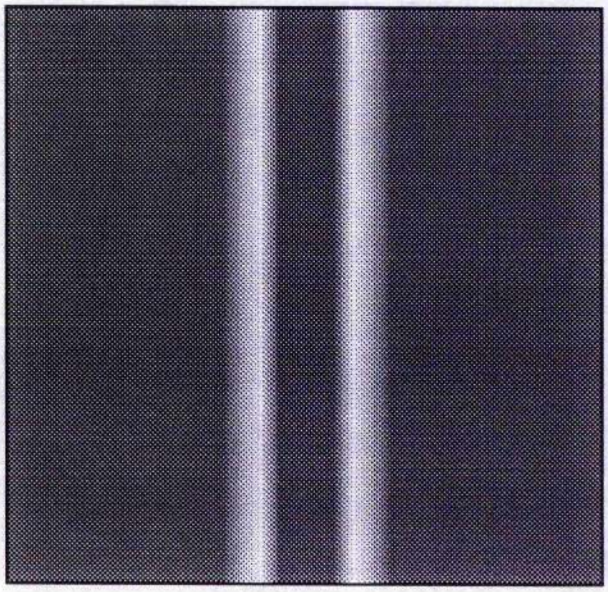

stimuli

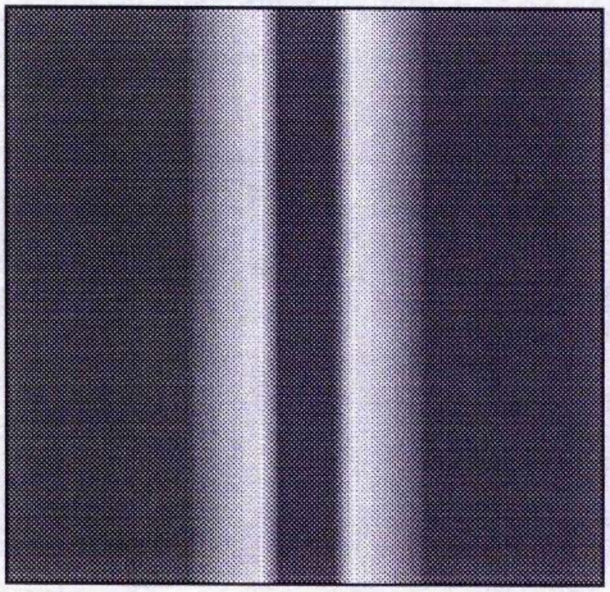

stimuli

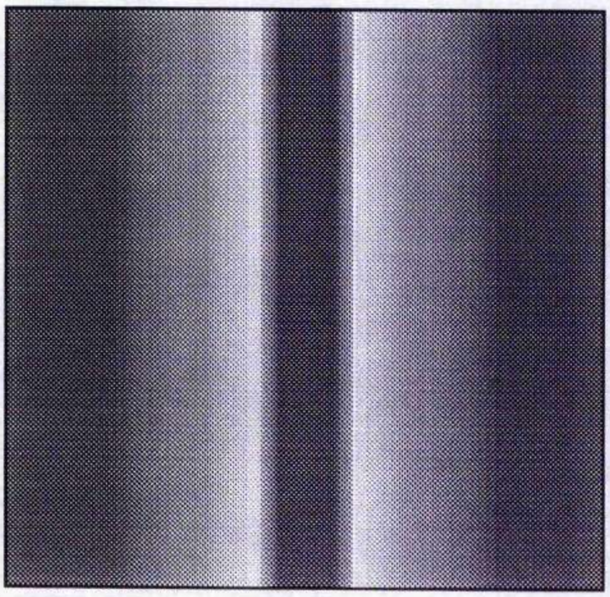

stimuli

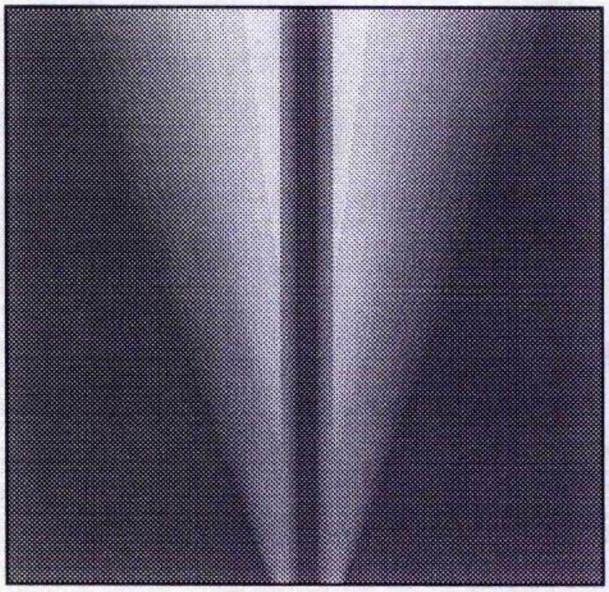

Chapter Five

Precision of spatial interval judgments during spatial perturbation

5.1: Overview

The purpose of the experiments reported in this section is to measure precision of judgments about the separation of adjacent luminance features in the face of spatial perturbations in the luminance profile of the stimuli, as a means of understanding how relative localisation of intensity changes in the image is performed. A number of spatial interval tasks were devised. Spatial interval thresholds were measured for targets consisting of truncated Gaussian luminance distributions (experiments 1-4) and for targets consisting of Gaussian bars (experiment 5). The first four experiments investigated the effect of increasing the skew of the contrast profile of the bars delimiting the spatial interval on the accuracy of interval discrimination. The effects of randomly varying either the degree of skew in the contrast profile, or the scale of the bars delimiting the interval on interval thresholds were determined. Thresholds were determined for briefly presented targets (20 msec) with random spatial modulation of part of the target.

The experiments where the contrast profile of the bars are randomly varied in some respect were designed to discover the information employed in encoding of spatial intervals by the visual system. The experiment involving brief exposure durations and longer exposure durations (20 msec and 500 msec) was an attempt to contrast the effects of randomising the degree of asymmetry of the contrast profile of the bars delimiting the spatial interval at short and longer exposure durations. There is evidence for larger filters being involved in coarse coding at brief exposure durations, and this may reflect in an increased difficulty for this interval task with the random variation in the degree of asymmetry of the contrast profiles of features delimiting the interval. Interval thresholds were also measured in the presence of a bar between the two bars delimiting the interval; this internal bar was scaled randomly in width, and in one condition, the skew of the outer bars was also varied randomly. Subjects had to compare the interval delimited by the inner edges of the two outer bars.

The data can be summarised as follows: interval discrimination was not affected by the degree of asymmetry of the contrast profile, as long as the inner section space constant was fixed. Randomising the space constant of the outer section of bars in an interval target (ie. outwith the luminance maxima) impaired interval discrimination slightly, elevating thresholds by an equal amount in the 20 msec and 500 msec exposure duration conditions, which was a surprising result. Whilst varying the outer section space constant affected interval thresholds only modestly, increasing the inner and outer space constants whilst keeping the separation of the zero-crossings in the intensity profile of the bars constant did increase interval thresholds more significantly. Randomly scaling these symmetric truncated Gaussian bars delimiting an interval increased thresholds by a small amount. Randomly scaling a third bar between the outer two target bars produced a small elevation of 1.5% in the Weber fraction for interval width, relative to baselines of 3-4%. Weber fractions of 6-8% were achieved when both the interior bar and the outer sections of the target bars were randomly scaled, which was unexpectedly low, since the target appearance from trial to trial was highly variable in this condition.

Spatial interval thresholds and interval ratio thresholds were determined in the face of perturbations in the orientation of the stimuli. Previous reports suggest that spatial frequency discrimination is as good between orthogonally oriented gratings; whether this is true for spatial intervals between bars, at intermediate orientations has not been clearly established, and it is possible that the grating result is artefactual, since it is known that subjects can use only one temporal interval in 2AFC spatial acuity tasks with equivalent accuracy to using both intervals. Manipulations were devised to force subjects to compare interval targets of different orientations. No clear dependence of threshold on orientation could be discovered in any of the experiments devised, leading to the conclusion that some simple spatial judgments are rotation-invariant.

The data from the 20 msec exposure condition in the second experiment, and modelling of the data from the third experiment (with random scaling of symmetric bars delimiting the target interval) make a MIRAGE operation unlikely, since Watt's (1987) data would suggest the involvement of some much larger filters at brief exposure durations, and extrema in these circumstances would no longer reliably encode the interval width in the face of the spatial perturbations imposed. However, if the MIRAGE operation is modulated by some higher-level control process, which can adaptively set the range of filter space constants involved in the discrimination, then this data does not present a problem for MIRAGE. This adaptive process might be triggered by having uninformative variance (ie. information which is not relevant to the task and must be treated as noise by the system) present at coarser spatial scales, as in several of the experiments reported here.

5.2: Introduction

There is still uncertainty about how the location of luminance changes in the retinal image are encoded. There is empirical support (Watt & Morgan, 1984, 1983a, 1983b) for the notion of spatial primitives. A spatial primitive for an edge is a way of characterising intensity changes in terms of information about the location, spatial extent and amplitude of the intensity changes. Spatial primitives are variants on the theme of a "local sign" (Lotze,1884), used by Hering (1899) to explain the precision of vernier acuity (from Boring, 1942). The original formulation of the local sign hypothesis involved averaging of information along the major axis of a vernier target to explain why vernier acuity should improve with line length. The dimension orthogonal to the major axis of a vernier target (the "orthoaxial" direction- Watt & Andrews, 1982) is also important in vernier judgments, and it has been suggested that some measure of central tendency of the intensity distribution in this dimension may be neurally represented (Badcock & Westheimer, 1985a; Watt, Morgan & Ward, 1983b; Westheimer & McKee, 1977a). Westheimer and McKee (1977a) reported that the perceived location of a narrow bright bar was controlled by the centroid of the retinal intensity distribution. They showed that subjects could discriminate the direction of offset of the centroid of a narrow, bright bar with an asymmetric contrast profile, with a precision which was equal to the best vernier acuity for symmetric bright bars (a precision of 3 sec arc was obtained); subjects were not aware of a luminance asymmetry in the asymmetric bar.

It is not a helpful observation that relative localisation is based on the retinal intensity distribution; this does not tell us a great deal about the processes which lead to perception. Marr and Hildreth (1980) suggested that early vision constructed a primitive representation of the image based on the responses of different scales of second differential filters. We can discover psychophysically what information is contained in this sort of representation by studying the pattern of visual system's information loss (the way in which the subject's error variance changes) as a function of stimulus parameters, and compare these with theoretical predictions about the nature of the representations in early vision.

5.2.1: Interval discrimination and vernier acuity

Whilst interval discrimination is a relative localisation task like vernier acuity, there are indications that it is dissimilar in some important ways. For fine abutting vernier targets, there is some limited dependence of vernier thresholds on the lengths of the lines; Westheimer and McKee (1977b) showed that vernier thresholds improve for target lengths up to 10 min arc. On the other hand, threshold for discriminating the separation of fine lines is not dependent on line length at all, at least for the separation of 3 min arc used in the experiment of Westheimer and McKee. Also, they showed that interval width thresholds are equal for targets composed pairs of positive contrast lines and pairs of negative contrast lines, sharp step edges and fine lines intermixed. However, a later study (Levi & Westheimer, 1987) showed that thresholds for intervals delimited by lines of different contrast polarities are higher at small separations (< 4-5 min arc) than for intervals delimited by lines of the same contrast polarity. Hence the discrimination of separation over small retinal distances and greater retinal distances may be subserved by different neural processes.

It is known that vernier thresholds depend on contrast. Watt and Morgan (1983b) showed that vernier acuity for blurred edges improves with contrast with an inverse square root relationship. Morgan and Regan (1987) demonstrated that contrast has no measurable effect on a spatial interval task beyond very low contrasts; they showed also that contrast randomisation of the features in the target only affects interval discrimination at very small separations (less than 2.5 min arc). Similar results concerning the lack of influence of contrast level on precision of judgments have been reported for spatial period discrimination (Caelli, Brettel, Rentschler & Hilz, 1983).

This implies a major difference between the neural mechanisms subserving vernier and spatial interval acuity. Morgan and Regan comment that it is very unlikely that spatial interval discrimination depends upon an independent extraction of local sign from the two bars, and explain the result in terms of a "coincidence detector" which receives input from two spatially separated, localised regions of the visual field. The spatial extent of these hypothetical regions is not known; the results of Morgan and Ward (1985) indicate that they may be quite small. Morgan and Ward reported that interval discrimination for two-line targets is not impaired by having randomly positioned flanking lines either side of the target lines. The results could only be explained by supposing that high frequency filters were responsible for encoding interval width; lower frequency filters predicted more interference with thresholds than was observed.

It is known that vernier acuity is degraded by increasing the stimulus blur (eg. Watt & Morgan, 1983b). There have been few if any studies relating interval width thresholds to blur of the features delimiting the interval. An extensive study of interval bisection at different blurs was conducted by Toet, Eekhout, Simons and Koenderink (1987). Their stimuli were presented at very low contrasts- ie. detection threshold- and consisted of blobs with two-dimensional Gaussian contrast profiles. Beyond small blurs of blob, thresholds for interval bisection were linearly related to the stimulus blur; ie. fractional thresholds were the same for all blurs regardless of the scale of the target (Weber's law). This result is analogous to the finding in spatial frequency discrimination experiments that a constant fractional difference in frequency can be discriminated across different frequencies (Campbell, Nachmias & Jukes, 1970).

There are some unusual aspects to interval discrimination which distinguish it further from vernier acuity. Vernier thresholds for sinewave gratings vary with spatial frequency (Bradley & Skottun, 1987), but a similar dependency on spatial frequency components of the stimuli is not observed with interval discrimination tasks. Burbeck (1988) showed that interval discrimination for pairs of bars is not affected by high spatial frequency modulation of the bars, or having a high frequency modulated bar paired with a lowpass Gaussian bar. A similar result was established by Toet and Koenderink (1988). They measured interval bisection and alignment thresholds for stimuli composed of three Gabor patches (the sinusoidal modulation of each patch was perpendicular to the main axis of the target). Threshold for these relative location tasks were entirely independent of the spatial frequency of the sinusoidal modulation over a 10-fold range (0.6-6.0 c/deg) and dependent only on the spatial extent of the Gaussian windowing function, replicating the results of a previous study by Toet, Eekhout, Simons and Koenderink (1987). Toet and Koenderink concluded that a mean patch location was being estimated from features derived from the entire intensity distribution of the Gabor patches.

One similarity between interval discrimination and vernier acuity is that both appear susceptible to the interaction between closely spaced contours. Levi and Westheimer (1987) reported that a line midway between two other lines changed the apparent width of a spatial interval, the change in the perceived width of the interval depending on the separation of the inner line and the outer lines; for separations of less than 3 min arc, the features were perceptually attracted (see section 3.2 of chapter 3). This was closely similar to an earlier study by Badcock and Westheimer (1985a), which looked at perceptual shifts produced by closely spaced features in a vernier acuity task.

There are models of visual processing that lay emphasis on the frequency amplitude spectrum. These models involve some computation on the outputs of local spatial frequency selective analysers (Wilson & Gelb, 1984; Carlson & Klopfenstein, 1985; Wilson, 1986). Wilson and Gelb demonstrated that data from a study of spatial period discrimination (Hirsch & Hylton, 1982) could be explained by their line-element model of spatial discrimination. The segmented nature of the function relating fractional difference thresholds to spatial frequency was used by Wilson as evidence for his line-element model, which was devised by analogy with line-element models in colour vision, which explain the well-known segmented nature of the hue discrimination function. On the other hand, others have failed to replicate the results of Hirsch and Hylton (Westheimer, 1984). The reliability of models like Wilson's depends on the fidelity of the local spatial frequency amplitude spectrum in signalling changes in the stimulus configuration related to performing the task. Random perturbation of the pattern's amplitude spectrum would be expected to disturb the operation of such models, by adding uninformative variance to the outputs of the local spatial frequency selective filters. This prediction was tested by Morgan and Ward (1985), in a study of spatial interval discrimination. Randomly perturbing outer flanking lines in a four bar comparison stimulus did not influence the precision of discrimination of separation of the inner two lines. These authors argued that their results provided strong constraints on models of discrimination which relied solely on information from the amplitude spectrum. Morgan and Ward suggested that the visual system might selectively use the output of relatively small scale filters to compute the interval separation. A similar conclusion was reached by Burbeck and Yap (1990) in a recent study. Watt (1987) has provided evidence for the existence of a coarse to fine analysis (ie. across spatial scale) in geometric tasks (such as assessing curvature, length or orientation discrimination). Burbeck and Yap sought to uncover an interaction between the exposure duration for an interval width target and the spatial frequency content of the stimuli. They measured separation discrimination thresholds for several different stimuli. One stimuli consisted of two target bars, each target bar having a pair of flanking bars (each target bar had one flanking bar inside and one outside the interval formed by the pair of target bars). An effect of varying exposure duration (ie. precision of interval judgments improved with increasing exposure duration) was found for this six bar target, but *not* for a two bar target. Note that these bars were separated by 2.9 degrees of visual arc. For the flanked bar targets, the inter-bar separations were all greater than 25 min arc, to prevent spatial interference between the bars. Additionally, they attenuated the medium and high spatial frequency content of the targets with a diffusion screen; with this screen in place, the exposure duration effect was greatly exaggerated for the six bar target. This they suggested was consistent with the "a priori selection of the high spatial frequency filters as the preferred source of information for this task" and not consistent with an automatic selection of filter scale, as in the original coarse to fine scheme of Watt (1987). They proposed a heuristic for scale selection; choose the scale of filter with the highest signal to noise ratio for task-related information; ie. "if the targets are presented in an uncluttered field, the strongest relevant signal comes from the larger spatial filters. If the observer has no prior knowledge of the stimulus or task, it may be that she will scan the filters from a coarse to fine scale because the coarser filters often give the strongest response, especially initially". Burbeck and Yap advocate a "position" model of relative localisation similar to Watt and Morgan (1984), which stands in opposition to models which use only information in the spatial frequency amplitude spectrum; a position model supposes that the relative location of several image features is determined by comparison of the spatial location of position tokens (ie. spatial primitives).

The present experiments were designed to investigate whether precision of interval discrimination could be impaired by spatial perturbations. Experiments 2-4 are related in some respects to Morgan and Ward's (1985) original study of the effects of spatial pertubations on interval discrimination. These experiments seek to determine the information necessary for interval judgments by extending the original study of Morgan and Ward to stimuli composed of features which have random variations in the contrast profile of each feature.

5.2.2 Interval discrimination and orientation

There are some visual judgments for which there is apparently no impairment of precision in the presence of an orientation difference between the discriminanda. For instance, Burbeck and Regan (1983) and Bradley and Skottun (1984) both reported that spatial frequency discrimination was equally good when the gratings to be discriminated were either orthogonally oriented or parallel with one another. The results of these grating frequency discrimination studies are not conclusive, since they both used only one spatial frequency in each session, and it is known that subjects can judge spatial intervals with normal precision when only one interval is present and feedback is given (Westheimer & McKee, 1977a). Using only one spatial frequency of grating in each session allows for the possibility that subjects may not compare the spatial frequencies across the different orientations, but rather just one of the orientations. For instance, subjects may monitor just the changes in spatial period of a vertical grating relative to some internal representation of a vertical standard- no comparison between the orthogonal stimuli is necessary to achieve asymptotically low performance, given the result of Westheimer and McKee, who obtained thresholds of 3 sec arc for a spatial interval discrimination with a 'learned' reference.

The results of Burbeck and Regan (1983) and Bradley and Skottun (1984) may be an artefact, since it is possible that the subject may rapidly "learn" an internal reference frequency for horizontal and vertical gratings independently, which an adaptive system like the brain can do without feedback. All that the subject has to do is to assess whether the grating at one of the orientations is of higher or lower spatial frequency than the mean spatial frequency of previous presentations. The experiment 5 here attempts to eliminate this potential artefact by: (a) randomising the reference and comparison orientations, so that the subject can not just develop one internal standard for (say) horizontally oriented gratings, (b) randomly interleaving different closely spaced reference interval widths, to prevent the subject basing the judgment on only one pattern.

5.2.3: Overview of experiments

The purpose of the experiments 1-4 in this section is to constrain candidates for spatial primitives which might be used in assessing interval width. Levi and Westheimer (1987) showed that the perceived width of the interval delimited by narrow, bright bars depends on the centroid of the retinal intensity distribution of the bars delimiting the interval. Morgan and Regan (1987) suggested that spatial primitives were not extracted to compare the relative location of the bars in a spatial interval target, since they observed no dependence of interval discrimination on target contrast. However, the result of Levi and Westheimer would imply that some sort of spatial primitive may determine perceived width of the interval; since they did not consider any candidates in a model of the neural response, they could not constrain the choice of primitive. Whilst the location of the centroid of the target bars in their experiment corresponded well with the perceived separation, other possible primitive features would show a similar variation in location; the zero-crossings in a small second order differential filter are also shifted, as are other features such as the extrema in this filter and the centroids of zero-bounded distributions in the filter response.

In the first experiment, targets are presented in which the locus of the centroid of the retinal intensity distribution is varied systematically, and thresholds determined as a function of the width of the outer section of the target. The luminance maxima are kept a constant distant apart, and the inner section of the luminance profile (between the two luminance maxima) of each bar is held at a constant width. The space constant of the Gaussian distribution which the outer section is based upon is varied systematically, to change the locus of the centroid of the retinal intensity distribution. If the centroid of the whole of the retinal intensity distribution of each bar delimiting the spatial interval is computed by the visual system, then a dependency of threshold on the space constant of the outer section will be observed. Weber's law implies that there is a constant fractional error in judging spatial separation, and has been demonstrated many times in spatial interval discrimination tasks (eg. Watt & Morgan, 1984; Hirsch & Hylton, 1982); thresholds should be proportional to the separation of the centroids of the intensity profile if the visual system is using a centroid measure in the intensity profile to assess interval width.

Further experiments were devised to investigate the encoding of interval width, which allowed for the possibility of subtler encoding strategies by the visual system. In a second experiment, the space constant of the outer section of each bar delimiting a spatial interval was randomised, and the effects on the precision of interval discrimination were measured. The third and fourth experiments followed up this basic approach to studying interval coding. Experiment three involved using stimuli with random scaling of the whole of the bars delimiting the interval, keeping the separation of the inner edges of the bars delimiting the interval constant (hence randomly varying the separation of the luminance maxima); experiment four involved randomly scaling a bar placed within the interval, loosely analogous to Levi and Westheimer's (1987) experiment, where they placed a line randomly perturbed in contrast between two target lines in a spatial interval task. In experiment four in the present study, the space constant of the outer section of each target bar as well as the space constant of the interior bar were randomly varied.

The aim of the fifth experiment was to determine whether previous reports that spatial period discrimination was insensitive to the presence of a 90 deg orientation difference between the discriminanda was due to an artefact, and whether this rotation-invariant visual judgment also operated at intermediate orientations in a spatial interval task. Interval thresholds were determined for stimuli consisting of two Gaussian bars, in the presence of perturbations in the orientation of the stimuli to be compared.

5.3: Method

5.3.1: Design

Spatial thresholds were determined for stimuli with differing spatial configurations, composed of either Gaussian bars or bars with asymmetric, truncated Gaussian contrast profiles.

5.3.2: Subjects

Main subjects used were MEW, GLC and IRP. Some data was collected from HAP. All wore their normal optical correction and none had astigmatism greater in size than 0.5 D.

5.3.3: Apparatus

As described in Methodology chapter.

5.3.4: Experiment 1: Interval thresholds for bars with skewed contrast profiles; stimuli

Patterns were composed of two bars. Each bar consisted of two halves. These halves were portions of Gaussian distributions, truncated at the inflexions in the distribution. This truncated Gaussian was shifted downwards to ensure that the luminance returned to the background level smoothly at the point at which the distribution was truncated. When symmetrical, this bar is very similar to a raised cosine bar (eg. Hines, 1976). The two halves of the bar were derived from Gaussian distributions which had (generally) different space constants.

The inner sections of these bars delimited a spatial interval equal in luminance to the background. The separation of the luminance maxima was held constant throughout one session, as was the separation of the inner bounds of each bar. Hence the location of the centre of mass of the intensity distribution of each bar could be shifted independently of the luminance maxima and inner zero-crossings in the luminance profile of each bar. A temporal 2AFC paradigm was employed; the subjects had to indicate which of two intervals presented was judged to be the wider.

Luminance profiles for these stimuli can be found in the appendix. The patterns presented were vertically oriented, and had a vertical extent of 45.0 min arc at 228 cm, or 90.0 min arc at 114 cm viewing distance. Reference interval separation, defined as the separation between the maxima in the intensity distribution of the stimulus, was either 5.0, 10.0 or 20.0 min arc. For the 5.0 and 10.0 min arc separations, the inner section always had a space constant of 1.0 min arc. Viewing distance for these patterns was 228 cm. For the peak-to-peak separation of 20.0 min arc, the width of the inner section was 2.0 min arc, and the viewing distance 114 cm. Width of the outer sections ranged from 1.0 to 45 min arc. Stimuli were visible for 500 msec, with a 500 msec ISI.

5.3.5: Experiment 2: random modulation of outer section space constant; stimuli

The stimuli used were identical to those above, but the space constant of the outer section of each bar delimiting the spatial interval was randomly and independently modulated by pseudorandomly generated scale factors (generated at run time), which ranged from 0.3 to 1.0. The scale factor was generated by a linear congruential method, which generates a uniform probability density function. Note that independent scale factors were generated for both the outermost sections of the two bars in the target. The inner section (ie. between the luminance maxima) was held constant throughout the experiment. Hence the luminance distribution between the luminance maxima remained constant throughout. The median reference separation of the luminance maxima was 10 min arc, with a range of 4 min arc. At this separation, Morgan and Regan (1987) found no effect of contrast randomisation of the target, and Morgan and Ward (1985) reported no effects of varying the spatial configuration on precision of interval judgments. The inner section width for all patterns was 2.0 min arc, and held constant throughout the experiment. Presentation time was either 500 msec or 20 msec with an ISI of 500 msec. Stimuli were viewed from a distance of 228 cm.

5.3.6: Experiment 3: random modulation of space constant of symmetrical bars- stimuli

Stimuli had the same luminance profile as in experiment 1, but the bars delimiting the interval were symmetrical. Two bars with truncated Gaussian contrast profiles delimited a spatial interval. The spatial interval between the inner edges of these bars (ie. the zero-crossings in their intensity profle) was kept a constant width.

The inner edges of these bars were separated by either 4.0 or 8.0 min arc. For the baseline stimuli, space constants of the bars delimiting the interval varied from 4.0 to 16.0 min arc; hence the separation of the luminance maxima of these baseline stimuli was not less than 12.0 min arc, and not greater than 42.0 min arc. A random bar width condition was also devised. The interval was defined as the separation of the inner edges of the bars, and this was kept constant; that is, the inner edges of each truncated Gaussian were kept in a constant location, whilst the space constant of the Gaussian from which the bar was constructed was varied randomly. This was done by generating a scale factor in the range 0.25-1 at run-time, by a linear congruential method. Data was collected from both the random bar width condition and the baseline condition where the width of the bars in the target was not varied. Exposure duration was 500 msec.

5.3.7: Experiment 4: random modulation of space constant of interior bar, and outer section of outer bars delimiting interval

Patterns used in this section were composed of three truncated Gaussian bars constructed as described for experiment 1 above. Interval separation is defined as the separation of the luminance maxima of the outer bars in the target, and this was varied between 12 and 40 min arc. Three conditions were tested; in the baseline condition (I), the bars had the same space constant throughout. In the second condition (II), the space constant of the interior bar was randomly modulated, by a pseudorandomly generated scale factor in the range 0.25-1.0. For the third condition (III), both the space constant of the inner symmetric bar and the space constants of the outer asymmetric bars were randomly modulated. These scale factors were independently generated for the outer sections and the inner bar at run-time, and could assume any value in the range 0.25-1.0. Hence the target in this condition (III) could have a highly variable appearance. Inner section space constant of the outer bars was held constant at 2.0 min arc. The interior bar was separated from the outer bars by not less than 3 min arc (2 min arc for the 12 min arc separation). For the baseline condition, the space constant of the interior bar was chosen as the median of the range over which the space constant was modulated in conditions II and III. Exposure duration was 500 msec.

5.3.8: Experiment 5: Interval discrimination for differently oriented stimuli

Patterns were composed of two Gaussian bars, of positive contrast polarity. Space constant of the bars was 1.0 min arc; separation was 16.0 min arc. Two distinct procedures were adopted: (1) randomising the reference orientation, to minimise the possibility of the subject basing the interval width judgment upon just one stimuli, (2) an alternative design involved holding the orientation difference constant in one session, with random selection of the reference separation. Reference separations were chosen in the range 8-16 min arc. Five independent measurements were made of interval threshold within this range, and the data pooled. The reference separation was randomly chosen for each trial. Stimuli were visible for 500 msec with an ISI of 500 msec.

5.4: Results

5.4.1: Experiment 1: interval discrimination for different asymmetries of luminance profile

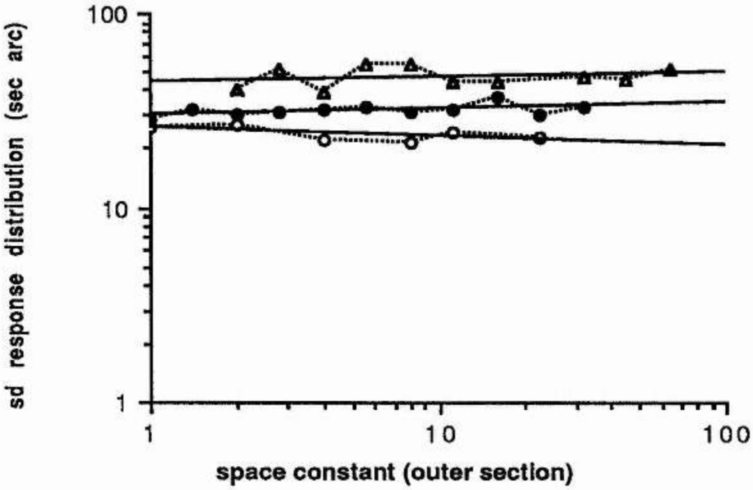

Figure 5.1: Data is shown for discrimination of interval width for three observers (IRP, GLC and MEW). Open circles represent data from the condition where the luminance maxima are separated by 5.0 min arc; filled circles are for the 10.0 min arc separation, and the open triangles represent data from the 20.0 min arc condition. Data from the 20.0 min arc condition is pooled across only two observers (GLC and IRP). There is no effect of the degree of asymmetry of the contrast profile on thresholds for this task. Error bars have been omitted for the sake of clarity.

Mean thresholds from all subjects were very similar for all conditions. These are presented above. There was no difference between the subjects in mean performance levels, and in the effect of the contrast profile skew of the bars on discrimination of the interval between the bars, and it therefore seemed justifiable to pool the data across subjects. Pooled data is shown above, with thresholds plotted as a function of the space constant of the outer section of the stimulus. The basic finding was that there is no dependency of precision of interval discrimination on the skew of the contrast profile, provided that the inner section of the bars remains constant throughout, and this is true for all the separations tested. Note that the separation of the centroids of the intensity profile increases with increasing space constant of the outer section, and if these features were localised, then a greater dependency of threshold on the outer section width would have been evident.

5.4.2: Experiment 2: interval discrimination whilst outer section space constant is varied

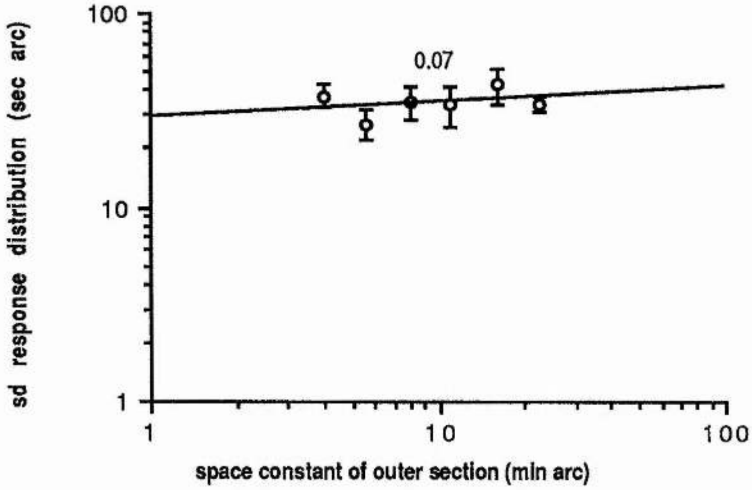

Figure 5.2: Above is shown pooled data from two subjects (IRP and MEW). Interval thresholds are shown for targets consisting of pairs of truncated, asymmetric Gaussian bars. The median separation was 10.0 min arc, and the exposure duration for the data presented above was 20 msec. Space constant of the inner section was 2.0 min arc. There is no clear effect of increasing the degree of asymmetry of the contrast profile on precision of discrimination; the power law function fitted to the data is nearly flat.

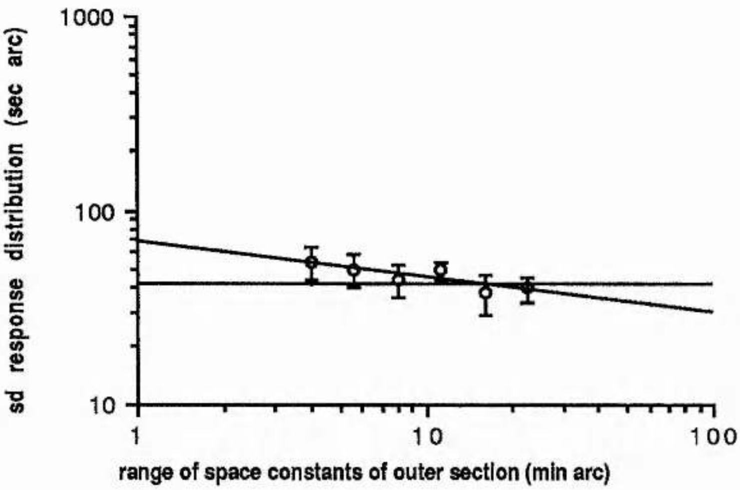

Figure 5.3: Interval thresholds are shown for pairs of truncated, asymmetric Gaussian bars with a median reference separation of 10.0 min arc, and an exposure duration of 20 msec, observer IRP. The baseline for this experiment is shown as a horizontal, which represents the average of estimates of separation thresholds with various different space constants for the outer section. The previous experiment indicated that this did not influence precision of the judgments, and this was basically true for the data from the 20 msec exposure duration condition (see previous graph). The function fitted to the data (by an unweighted least squares fit) is a power law function, which has a negative exponent (-0.18).

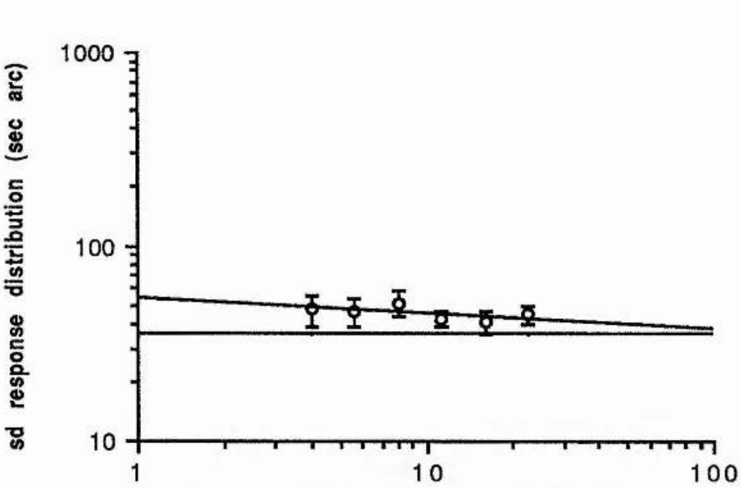

Figure 5.4: Interval thresholds are shown for an exposure duration of 20 msec, observer MEW. Baseline (with no spatial perturbation) shown as a horizontal line. The function fitted to the data is a power law function, fitted via an unweighted least squares fit. This function has a negative exponent (-0.07).

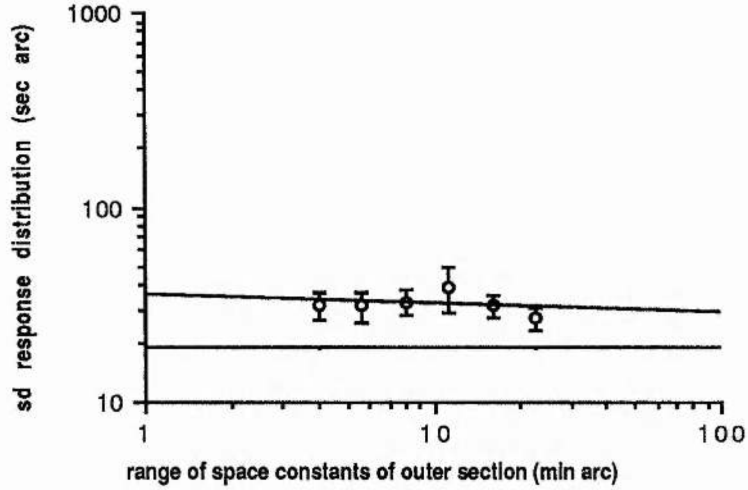

Figure 5.5: Interval thresholds are shown for an exposure duration of 500 msec, observer MEW. Baseline (with no spatial perturbation) is shown as the lower line. The function fitted to the data is a power law function with a small but negative exponent (-0.04).

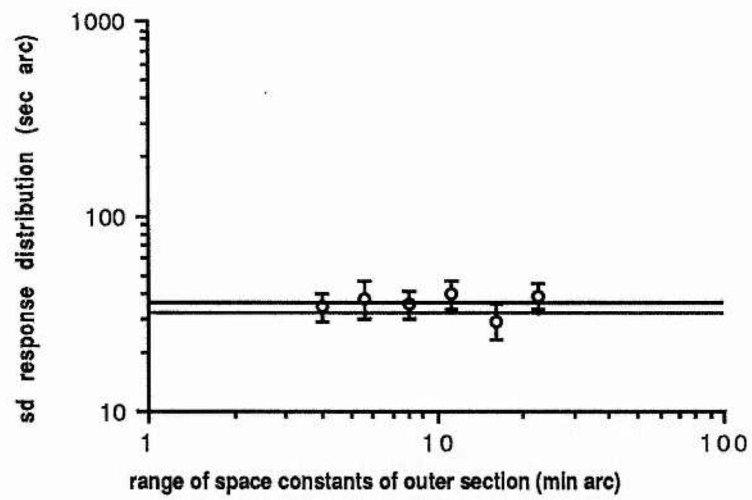

Figure 5.6: Interval thresholds are shown for an exposure duration of 500 msec, observer HAP. Baseline (no spatial perturbation of the target) is shown as the lower line. The upper line is a power law function- its exponent is very close to zero (-0.0004) but still negative.

Data is presented above for the 20 msec and 500 msec exposure durations. The data indicate that in most cases, increasing the range over which the width of the outer section of each bar is increased leads to an *improvement in precision*, reflected in the negative power law exponents for the functions fitted to the data. This might appear paradoxical, but the range over which the outer section is varied begins at increasingly greater separations; the minimum width of each bar is the width of the inner section plus 30% of the maximum outer section width. Also, there need be no monotonically increasing relationship between the range over which the space constant is varied, and the range over which separation of features which are used to represent the separation of the bars vary (if such a process is even involved) .

There is no clear interaction between the range over which the outer section's space constant was randomised and the exposure duration; ie. thresholds were not differentially affected by making the target vary in overall width at the 20 and 500 msec. The mean threshold elevation is quite similar in each of the exposure duration conditions; randomly varying the degree of skew of the luminance profile of the stimulus did not make the task any more difficult at a 20 msec exposure duration, than if the stimulus was visible for longer. If anything, the randomisation of the spatial configuration made the task more difficult (in terms of the threshold elevation measured) at the longer exposure duration (at least for observer MEW). This was a surprising result. Thresholds for the 20 msec exposure duration are still in the order of 30-40 sec arc, which indicate that the judgment cannot likely be performed simply as a resolution judgment (paired versus unpaired), since the psychophysical procedure used would not usually present any stimuli which were more than 3 standard deviations of the response distribution below the reference separation (median reference separation being 10 min arc) ; in the worst case, a discrimination would be made between separations of, say 6.5 and 8 min arc, and the 6.5 min arc separation is still resolvable.

5.4.3: Experiment 3: interval discrimination for symmetric targets where the bars are randomly scaled in width

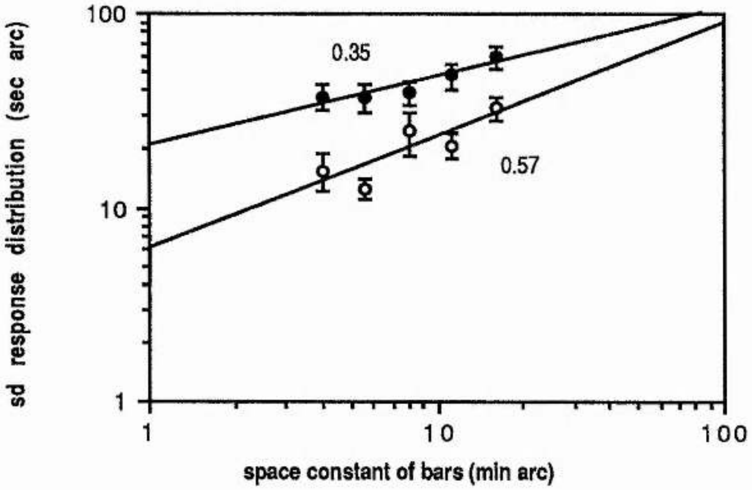

Figure 5.7: Data from discrimination of separation of pairs of symmetric, truncated Gaussian bars as a function of the space constant of the bars is shown above for two different reference separations- open circles represent the 4.0 min arc separation of the bars' inner zero-crossings in the intensity profile, and the filled circles represent the 8.0 min arc separation of the inner zero-crossings. There is no random perturbation of the stimulus configuration for the data presented above. The data has been pooled across MEW and IRP for the 4.0 min arc separation, MEW,IRP and GLC for the 8.0 separation, and is based on up to two determinations per subject, each involving 160 trials. The continuous line represents the best-fitting power law function, fitted by an unweighted least-squares fit. As the width of the bars increases, there is a gradual increase in the threshold for discriminating the width of the interval between them. Width of the bars is twice their space constant, since they are truncated at the inflexions in the intensity profile. The effects of this manipulation are less marked at the greater interval width, reflected in the decreased power law exponent.

The separation of the inner zero-crossings in the intensity profile of the bars is kept constant, and the separation of the luminance maxima increases as a linear function of the space constant of the bars. However, the luminance gradient at the edges of the bars also decreases with increasing space constant, and since this is directly related to the edge contrast, this may explain why the task gives a function with a moderate power law exponent (0.35-0.57). Watt and Morgan (1983b) obtained power law exponents of close to -0.5 for the effects of contrast on vernier acuity. However, this is not conclusive, and to show that the subjects are localising only the inner edges of the bar (which is implied heavily by the results of the first two experiments), it is necessary to look at data from the condition where the size of the bars is being randomly scaled. One significant aspect of the data in the above graph is that varying the space constants has less effect for the wider separation. If the system is locating some features in the neural response distribution whose separation increases less slowly than the separation of the luminance maxima, then a power law exponent of less than 1 will be found for this task. If we increase the reference separation, then there will be less effect of increasing the space constant on precision of this judgment (ie. the power law exponent will go down) since the same rate of increase in the separation of the localised features is being added to a larger baseline separation. If we make the separation very large relative to the change in position of the localised features as we vary the width of the bars, then it is clear that in the limit the function relating the width of bars to the thresholds for their separation will be nearly flat, ie have a power law exponent of close to 0.0. This explains why the power law exponent should be less for the wider reference separation (0.35 for the 8 min arc separation, and 0.57 for the 4 min arc separation).

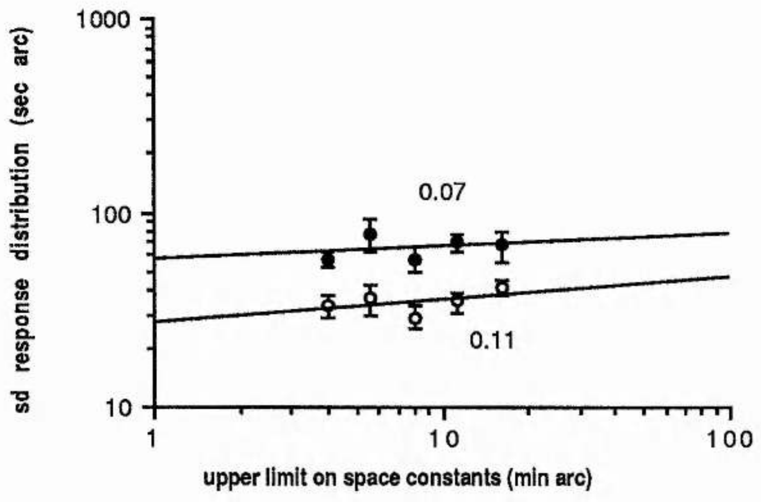

Figure 5.8: Data is shown for discrimination of the separation of two bars, which consist of symmetric truncated Gaussians, whose space constant is randomly varied across a range $0.25\ s < s_{random} < s$, where s is the upper limit on the range of space constants plotted above.. Open circles represent data from the condition where the inner zero-crossings in the intensity profile are separated by 4.0 min arc and filled circles are from the condition where the inner zero-crossings are separated by 8.0 min arc. Data shown is pooled across two or three subjects (IRP and MEW for the 4.0 min arc condition, and MEW, GLC and IRP for the 8.0 min arc condition). Whilst the appearance of the target on any one trial is highly variable, subjects can still perform the task with a fair degree of precision.

For the stimuli with the broadest space constant of outer section, in the random condition the outer section space constant varied from 4-16 min arc; the sd of the variation in the separation of the luminance maxima (and the centroids of the intensity distribution) is 6.9 min arc (the range divided by $\sqrt{12}$, which is the standard deviation of a uniform distribution). Since this additional variance would increase thresholds by a very large amount, these features are not relevant to perception of the separation of two bars. The power law exponents are much less than in the condition where the stimulus was not randomly varied in appearance; this is to be expected if the feature to be localised is not linearly related to the space constant of the inner section of the bars, and randomising the width of the bars over a large range may produce only a relatively small change in the position of the features being localised at the inner edges of the bars. The power law functions are used as a convenient approximation to the data, but no theoretical basis is claimed for them in these cases.

5.4.4: Experiment 4: interval discrimination in presence of randomly scaled bar in middle of interval, and random scaling of outer section of outer bars

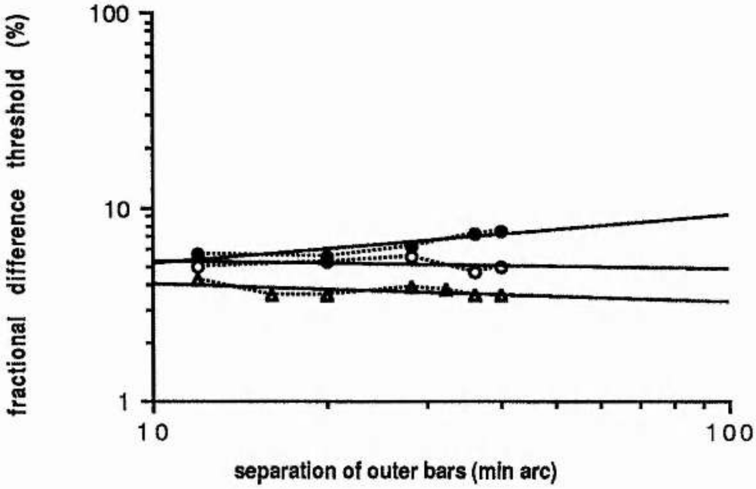

Figure 5.9: Data is shown above for one subject (IRP) for discriminating the separation of two bars with a third bar placed between. There are three conditions represented here; the open triangles are the baseline with no spatial perturbation, the open circles are the condition where the width of the interior bar is randomly varied, and the filled circles are the condition where the widths of both the outer section of the outer bars and the innermost bar are perturbed. Error bars have been omitted in the interests of clarity.

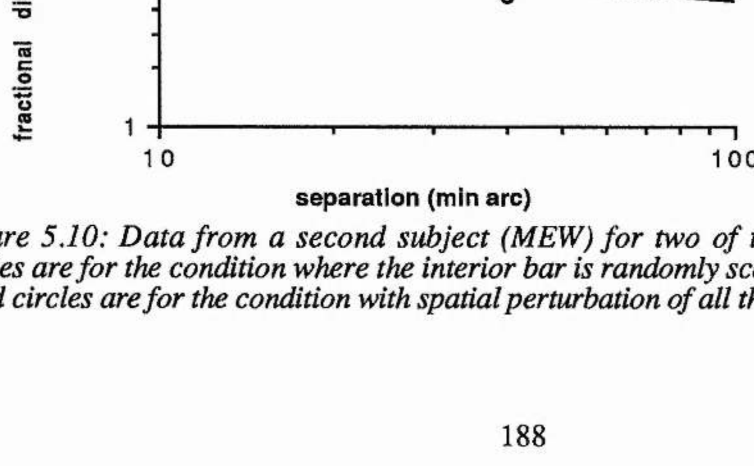

Figure 5.10: Data from a second subject (MEW) for two of the conditions; open circles are for the condition where the interior bar is randomly scaled in width, and the filled circles are for the condition with spatial perturbation of all three bars.

Thresholds for spatial interval discrimination for three bar targets with an irrelevant bar between the outer two asymmetric target bars are shown above. Expressing the thresholds as a Weber fraction indicate that precision is largely proportional to the separation of the inner edges of the outer bars (ie. the target separation), demonstrating a characteristic Weber's law like behaviour. The effect of adding a randomly scaled interior bar is to raise thresholds marginally for the interval discrimination (by about 1.5% on a baseline of 3-4% for subject IRP). The baseline for this task consisted of three bars, the innermost of the three having a width which was in the middle of the range over which the space constant was randomly varied in the other two conditions. Randomly varying the width of this middle bar elevated thresholds slightly for observer IRP; however, it is plain that the visual system can perform this task without the necessity to add together the widths of the adjacent spatial intervals in the target, whose combined width is subject to a high variability by this manipulation- the visual system can apparently independently extract information about the separation of the inner sections of the outer bars without necessity for an intermediate computation (ie. it can ignore the middle bar). Another possibility is that it measures distance from the inner edges of the outer bars to the centre of the central bar, and takes the sum of these distances. This possibility cannot be excluded. Randomising the space constant of the outer section of the target bars, in addition to the space constant of the inner symmetric bar increases thresholds a little further, but the Weber fractions are still in the region of 6-8%, indicating that despite wide fluctuations in the appearance of the stimuli to be compared, discrimination performance is still very good.

5.4.5: Experiment 5: interval discrimination during orientation perturbation

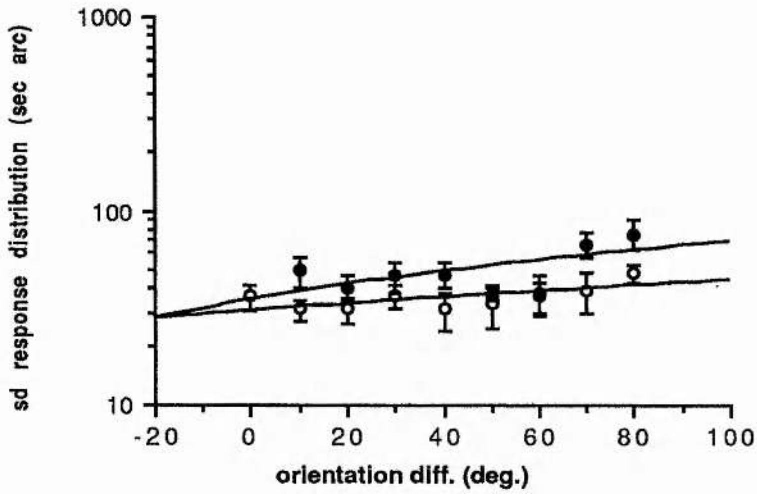

Figure 5.11: Thresholds for interval width judgments are illustrated for the condition where the reference orientation is chosen randomly. Data is shown for two subjects (IRP; open circles, GLC; filled circles). Interval width was 16.0 min arc, and the interval was delimited by two Gaussian bars with space constants of 1.0 min arc. There is a slight upward trend in the data, indicating that as the orientation difference between the discriminanda was increased, so precision on the task decreased slightly. The continous lines are straight-lines fitted to the data by an unweighted least-squares fit.

Thresholds for interval discrimination for targets with a random reference orientation are presented above, as a function of the orientation difference between the discriminanda.There is some small effect of introducing an orientation difference, but this is likely to be due to an artefact. The horizontal-vertical illusion is well known and indicates that subjects will perceive vertical intervals as narrower than horizontal intervals. Some of the pairs of stimuli which differed in orientation by 90 deg. would have been presented with one horizontal and the other vertical, whilst others would have been presented at intermediate reference orientations (although with the same 90 deg orientation difference). The pairs which were horizontal-vertical would introduce the horizontal-vertical illusion, but this would not happen for two pairs oriented, for instance, at 45 and 135 deg. This additional variance in the perceptual width of the intervals when there was a large orientation difference between the discriminanda could explain part of the threshold elevation. Consequently, thresholds were determined in the situation where any illusion of this sort would manifest simply as a shift in the mean of the psychometric function and not a change in its slope, by using a fixed slope of reference orientation, but randomly varying the reference separation, to avoid the subjects basing the judgment only on a single interval.

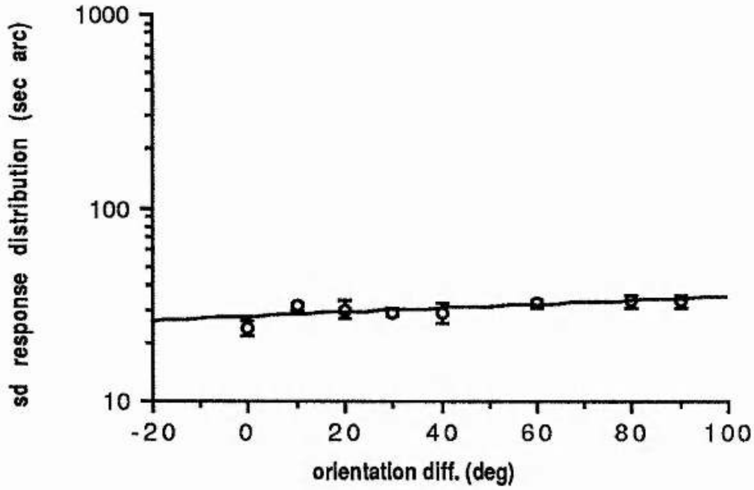

Figure 5.12: Data from one subject (IRP) for the condition where the reference separation is chosen randomly in the range 8-16 min arc. Median interval width was 12 min arc, and the interval was delimited by two Gaussian bars with space constants of 1.0 min arc. There is a very slight upward trend in the data, indicating that as the orientation difference between the discriminanda was increased, so precision on the task decreased slightly. Each data point is the arithmetic mean of five determinations, and is based on roughly 700 trials. The continuous line is a regression line fitted to the data.

Thresholds for interval discrimination with randomly chosen reference separation are shown above. The data is a replication of the above experiment under slightly different conditions, and data is reported from only one subject, although the results were confirmed with another subject. Since this was a replication of a previous experiment, data from only one subject is presented. There is some indication that having both stimuli vertically oriented leads to lower thresholds, but this may have been a practice effect, since the subject was highly familiar with vertical stimuli in an interval task and had less acquaintance with comparison of stimuli differing in orientation, so on balance the difference may not be genuine. In comparison to the previous experiment, only one orientation difference was tested in any one session; however, the reference separation was randomly chosen over an 8 min arc range as mentioned above.

5.5: Discussion

Precision of discrimination is not influenced by the degree of asymmetry of the contrast profile of the bars delimiting the interval, provided the space constant of the inner section is kept constant. Data collected during the course of the second experiment suggested that this was also true at 20 msec exposure durations. It was possible to influence precision of discrimination by altering the overall width of the bars delimiting the interval only if this involved a change in the space constant of the inner section of the bars (experiment 3).

Subjects are equally good at assessing the width of the interval between two luminance features regardless of whether the delimiters of the interval are phenomenal 'bars' or 'edges' (for the stimuli used in the first two experiments, if the asymmetry of the contrast profile is great, then the interval is seen as delimited by two edges) and show no measurable decline for features intermediate between these extremes. Westheimer and McKee (1977a) showed that this interval discrimination is equally precise for lines or edges with a 3 min arc separation and very sharply defined features. Randomising the spatial configuration of the target from trial to trial does have some effect; in the present study, Weber fractions for interval discrimination were increased by at most about 3%. As observed by Morgan and Ward (1985), manipulations like this lead to a gross decorrelation of the frequency amplitude spectra of the targets to be compared from trial to trial, with the small cue to the change in width submerged in the large trial to trial variations in energy at different spatial frequencies. Not all of the information contained in the spatial frequency amplitude spectrum need be used in judging interval width, and the results of the experiment here cannot entirely exclude the possibility of a localised frequency analysis.

One feature of the results with the randomly modulated stimulus configuration is that there is no discernible dependency of the accuracy of discrimination on the range over which the space constant of the outer section is varied. Discrimination is remarkably precise even when the target's overall width is randomly varied from trial to trial over a range up to 30 min arc and the stimulus is visible for only 20 msec- eg. fractional difference thresholds of as low as 5-6% were obtained reliably under these conditions. Also performance is remarkably good even with a third bar of variable width interposed between the two bars delimiting the target interval which randomly scaled in width over a wide range. Spatial interval discrimination displays a remarkable flexibility in the face of a variety of manipulations, which line-element models such as Wilson's (1986) would not predict at all without major refurbishments.

Calculating the locus of the centroid of the retinal intensity distribution of each bar indicates that the separation of the centroids of the intensity profiles of the bars varies over quite a broad range for the stimuli used in the experiment, yet there is no detectable effect of increasing the skewness of the luminance profile on thresholds for interval width. The results of plotting the relative locations of centroids in the stimulus intensity profile for the stimuli employed in experiment 1 can be seen below.

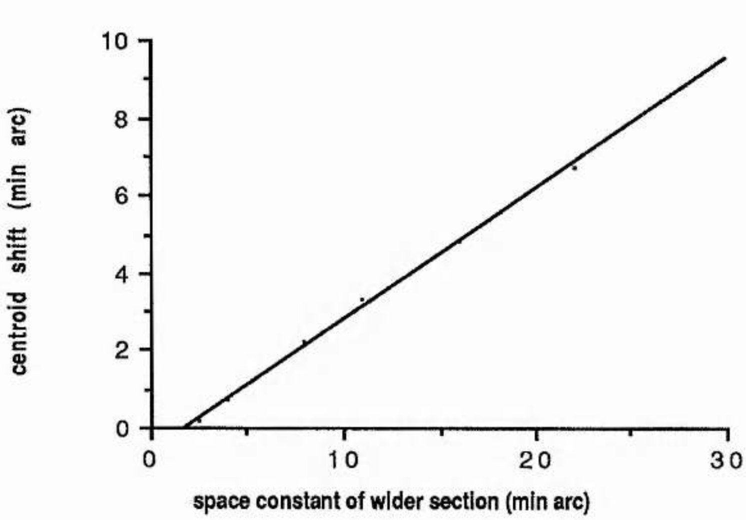

Figure 5.13: If the system measured the distance between the centroids, then we might expect a Weber's law relationship to hold between the centroid separation and threshold. Above is a graph of the function relating the location of the centroid of the asymmetric bars used in experiment 1 to the space constant of the wider portion of the bar, relative to the centroid location for a symmetric bar. The inner section space constant is taken as 2.0 min arc. The centroids of the bar pairs are 12 min arc further apart for the broadest bar widths employed relative to the symmetric stimulus.

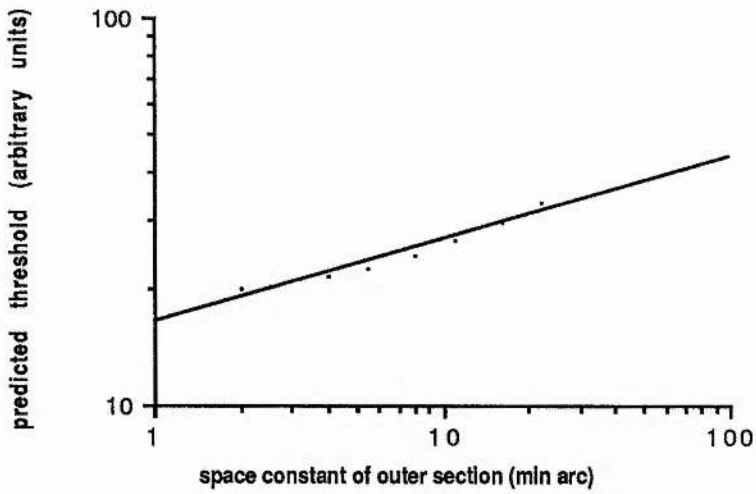

Figure 5.14: The above shows a theoretical relationship between the space constant of the outer section of each bar and a measure of separation of the centroids which ought to be proportional to threshold if a Weber's law relationship between centroid separation and threshold holds.

Since the above is a plot on log-log axes, a scaling factor (the constant of proportionality between centroid separation and threshold) will merely result in a vertical shift of the function. If interval thresholds increase in proportion to the separation of the centroids in the stimulus intensity distribution, then the data should follow a function which is well-fitted by a power law function with an exponent of 0.21, the exponent for the above plot. However, the data function relating thresholds to the space constant of the outer section was flat (see figure 5.1); that is, varying the outer profile of the bars did not affect thresholds for discrimination of their separation, leading to the conclusion that the centroid of the stimulus intensity profile is not represented by the visual system for these sorts of separation tasks.

This indicates that if the centroid of the intensity distribution is computed by the visual system, it can only happen over a relatively restricted area and not over the entire feature. The results of the third experiment (where symmetric bars delimiting the interval were randomly scaled in width, with the separation of their inner zero-crossings being held constant) eliminate the possibility that luminance maxima are involved in interval width perception. The standard deviation of the variation in location of the luminance maxima was 6.9 min arc, but a tiny elevation in interval threshold is seen (about 10 sec arc) in the face of this manipulation.

Consider for a moment the data from the condition in experiment 3 with no spatial perturbation. The bars were symmetric, and interval thresholds were measured as a function of the space constant of the bars, for a given separation of the inner zero-crossings in the intensity profile of the bars. Precision of interval discrimination declined as the space constant of the bars was increased. The results of Morgan and Regan (1987), which show that contrast has no effect on interval discrimination beyond very low contrasts, indicates that the decline in precision observed in the present experiment 3 cannot be attributed to the decrease in luminance gradient at the zero-crossings in the intensity profile as the space constant of the bars is varied, unless the bars are no longer treated as bars and are treated as edges.

5.5.1: Modelling

Several of the experiments described above involved addition of width jitter in the bars delimiting the interval, which was hoped to have the effect of introducing additional variance into the internal representation of the stimulus. This additional source of variance in the representation would be expected to decrease precision; unfortunately, some of the manipulations did not noticeably impair the precision of the judgments- indicating that the judgment was limited solely by the internal error in representing interval width. If we can find features which give roughly the same pattern of information loss shown by the human observer as we vary the stimulus noise, then we may have found the features which the visual system uses to assess the separation of the bars.For example, it was discounted above immediately that the system localised either luminance maxima or centroids of the intensity distribution- these features would add too much error to the precision of judgment and give thresholds well over 7 min arc.

The data from these jitter experiments can be modelled by a simple model (eg. Hess & Watt, 1990), which involves linear addition of independent sources of variance.

$$t = k. \sqrt{s_e^2 + s_i^2}$$

where t = psychophysical threshold

k = constant

s_e = externally applied error in stimulus property

s_i = internal error in representing stimulus property.

Data from experiment 2 and 3 is modelled using the above expression. Stimuli for this experiment were composed of truncated asymmetric Gaussian bars with an inner section space constant of 2.0 min arc. Peak to peak separation for the modelling was taken as 12 min arc. The outer section of the stimuli had space constants of 1.2 to 22.6 min arc.

Predictions were computed by: (1) generating sets of predictions for 25 different stimuli within the range of outer section space constants. For instance, if the upper bound on the outer section space constant was 4.0 min arc, given that the scale factor in experiment 2 was in the range 0.3-1.0, then the lower bound on the outer section space constant is 1.2 min arc. Predictions would be computed from models based on different sorts of spatial primitives for 25 stimuli with outer section space constants from 1.2-4 min arc. The predicted perceived location of the inner edges delimiting the interval were computed; since the contrast profile of each bar was varied independently, the variance in the internal representation of interval width caused by randomising the contrast profile of the outer section of the bars delimiting the target would be the sum of the variances in predicted location of each of the perceived 'edges' of the interval (since the total variance of a number of independent random variables is additive). This procedure was repeated for each range of outer section widths. Hence 150 sets of predictions would be generated for the range 1.2-22.6 min arc, given there were six sub-ranges contained in this range.

Modelling of the data from experiment 3 was performed identically; the variance in the predicted location of the 'edges' of the spatial interval was computed, and then the total variance in the interval width was taken as the sum of these variances.

(2) The variance of predictions about the representation of interval width was then added to the variance of the unperturbed interval width judgments, to give the estimated total variance in judgments about width of the interval. The square root of this variance is a prediction of the spatial threshold for the task. Some of the predictions are plotted below.

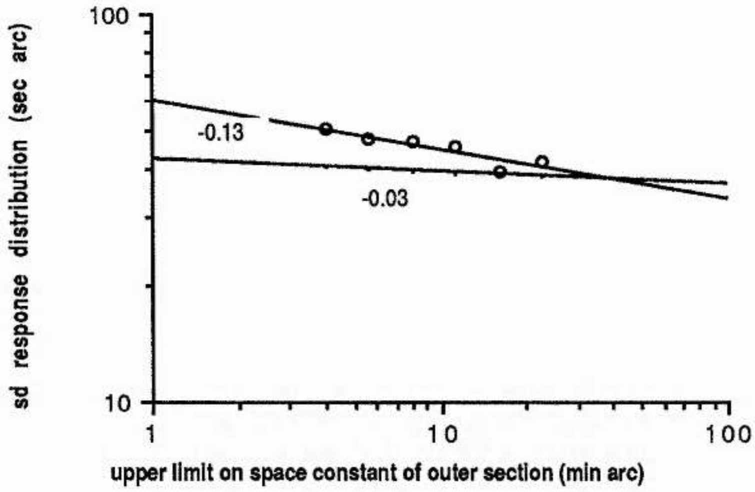

Figure 5.15: The open circles on the graph above represent pooled data from the 20 msec condition (pooled across MEW & IRP). The functions shown are power law fits to the empirical and simulated data. The lower function is simulated data generated by supposing that the visual system is localising zero-crossings in the output of a LoG filter with a space constant of 1.0 min arc. This gives a function which has a power law exponent of -0.03.

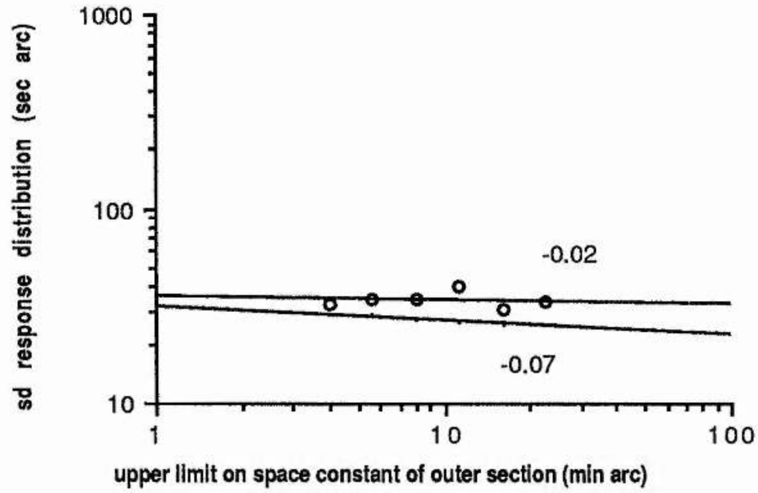

Figure 5.16: Data and predictions for the 500 msec exposure duration (data is pooled across MEW & HAP); the lower function represents predictions made by supposing that the visual system localises the zero-crossings in the output of a LoG filter with a space constant of 1.0 min arc. Note that the power law exponent has changed for these predictions compared with the set above, since the baseline variance (ie. the error variance of the subject for judgments about stimuli with no spatial perturbation) is lower for the longer duration, and the variance introduced externally has a proportionally greater effect.

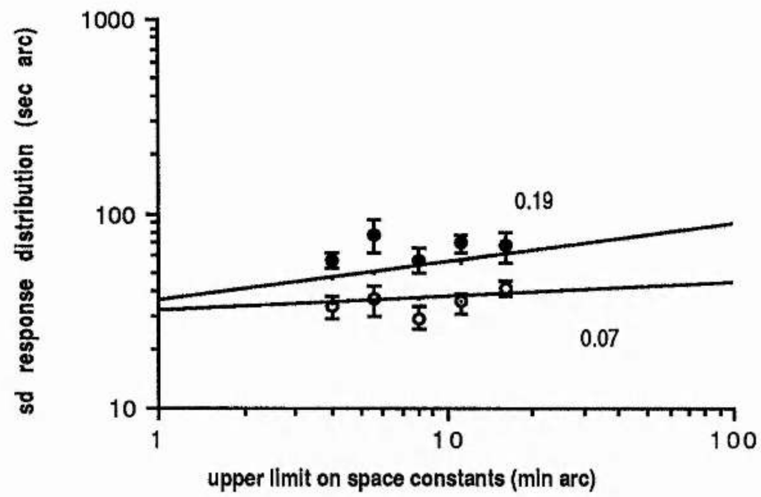

Figure 5.17: Comparison of predictions and data from third experiment; the above graph represents predictions of a model supposing linear addition of variances, due to internal and external noise. Pooled data is shown as in the results section, with the open circles representing data with the inner zero-crossings in the intensity profile of the stimulus separated by 4 min arc, the filled circles representing data from the 8 min arc separation.

The functions illustrated are power law functions fitted to simulated data. The simulated data was derived by supposing that the visual system is localising extrema in the output of a LoG filter with space constant of 0.35 min arc (the smallest filter in MIRAGE), and that the variance introduced by localising this spatial primitive adds linearly to the internal variance in registering interval width. This provides a reasonably good account of the data. If any larger filters are included, then the power law exponents become very much larger.

In the predictions displayed in figure 5.15, the internal error was estimated as the square of the threshold for the largest width of bar used (ie. the extremity of the range), and perhaps it would have been more appropriate to take an average of the thresholds from the range of jittered widths. If this is done, then the results can be seen below.

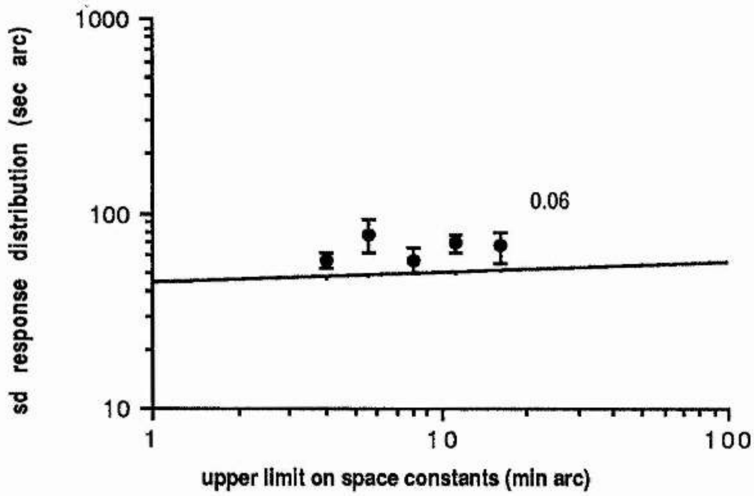

Figure 5.18 This is the same data set as the previous graph, but the simulated data has been generated by taking the internal error as the average threshold for the range over which the width of the bar was jittered. This gives a function with a power law exponent which is much reduced (0.19 has become 0.06), and now very close to that possessed by the original function fitted to the data (0.06 compared with 0.07), but the function now lies below the empirical data slightly.

5.5.2: Discussion of results of modelling

Modelling of data from the second experiment in this chapter, where the space constants of the outer sections of bars delimiting a spatial interval are randomly varied, suggest the data is reasonably well explained by supposing the system localises zero-crossings in a filter with space constant of 1.0 min arc (roughly the filter size used by Toet, Smits, Nienhuis & Koenderink, 1988 to explain their results). This gives a function with a small but negative power law exponent; the disparity between the data and predictions of this scale of filter for the 20 msec data may arise from the diminished visibility of some of the narrower bar pairs presented at this very brief exposure duration, which may have given artefactually increased thresholds for these data points (and hence given a higher power law exponent).

On balance, this experiment was not particularly conclusive, given that the slopes of the functions were so close to zero, and the results do not constrain too tightly the choice of primitives for interval width. When an experiment returns a null effect, there is always a high risk of a type II error, ie. to wrongly accept the null hypothesis; attempting to 'prove the null hypothesis' is a very bad experimental procedure. This was not the intention behind the second experiment, as the results were a complete surprise.

However, modelling of the data from the third experiment, where the symmetrical bars delimiting the interval were randomly scaled in width suggests that precision of judgments in this experiment can *only* be explained by supposing that features such as extrema in the output of a very small filter with a space constant in the region of 0.5 min arc; no larger filters can be involved on the basis of this data.

The width of an interval is clearly computed quite flexibly, using only sections of the neural response distribution which correspond quite closely to the perceptual 'edges' of the interval. The results are entirely inconsistent with Wilson's (1986) line element model for spatial discrimination, which uses only the energy in spatial frequency detectors to decide whether patterns are the same or different. This conclusion was reached by Morgan and Ward (1985) whose experiment was devised to test the intuition that fine spatial interval discrimination is possible for a target with large changes from trial to trial in the spatial frequency amplitude spectrum. The present experiment 2 was based loosely on the Morgan and Ward experiment, and it provides essentially a confirmation of their result. The model of interference with spatial interval discrimination presented here is certainly over-simplified; it cannot explain why interval discrimination should be independent of contrast (Morgan & Regan, 1987). In the present experiments, the subjects were specifically asked to attend only to the interval and ignore the bars as best they could. The results indicate that they can do this quite competently even for very brief exposures. These data and simulations seem initially inconsistent with Watt's (1987) coarse-to-fine scheme for representing images. Watt suggested that judgments such as length discrimination (conceivably another sort of interval discrimination) reflect the properties of the largest scale of filter active for a given stimulus exposure duration. The present data indicate that spatial interval discrimination is still possible with high precision at very brief exposure durations even in the face of irrelevant changes in the spatial configuration of the target which would introduce large variance in the output of low spatial frequency filters, indicating that high frequency filters are still used exclusively in this situation. This result in itself was quite surprising, since the initial assumption behind comparing the two exposure durations in experiment 2 was that randomising the target's overall width would be highly deleterious for brief presentations. It was not apparent before the experiment that this result could be possible, since it is somewhat counterintuitive. MIRAGE cannot explain why discrimination should still be so accurate under the conditions in experiment 2, where exposure duration is 20 msec, unless the range of filters active is similar to those active in processing the 500 msec duration. If any larger filters are active, then the location of the extrema (or centroids) in the MIRAGE output is shifted outwards greatly.

5.5.3: Orientation perturbations and interval width discrimination

The lack of dependence of threshold on orientation difference between the patterns to be discriminated when care was taken to avoid possible artefacts due to use of only one stimuli to make the judgment was a surprising result. Since the orientation of the reference and comparison target were varied pseudo-randomly, subjects were required to compare the perceived width of the two spatial intervals, rather than relying on just one of the intervals and comparing this to some internal 'standard' which subjects can do with an equivalent accuracy to judgments made in the more usual 2-AFC paradigm (Westheimer & McKee, 1977a).

5.6: Summary and conclusions

A variety of spatial interval discrimination tasks were performed. Interval thresholds for bars with asymmetric orthoaxial contrast profiles were not affected by the degree of asymmetry of the contrast profile, as long as the inner section space constant was fixed. This suggests that the separation of the centroids of the intensity distribution is not necessary information for judging the separation of two luminance features. Randomising the outer section space constant of bars delimiting a spatial interval (equivalent to randomly varying the degree of asymmetry of the luminance profile) over a wide range had little effect on interval discrimination, even when exposure duration was as brief as 20 msec; the effect of this manipulation was to increase thresholds by at most 10 sec arc. A similar increase in threshold was seen when the interval was delimited by symmetric bars, whose space constant was randomly varied (the separation of the inner edges of each bar being kept constant). MIRAGE cannot account for the data from the 20 msec condition without assuming that the range of filters active is the same as in the 500 msec exposure duration. Additionally, modelling of the data from the third experiment (with symmetric bars randomly scaled in width) suggests that performance is too good for medium spatial frequency filters to be involved, and is not consistent with a MIRAGE operation; however, the degradation in interval discrimination is consistent with the extraction of information about the location of features such as extrema in the response of a high spatial frequency filter (eg. a LoG with a space constant in the region of 0.5 min arc).

Interval thresholds were measured for targets with a third bar between the outer two target bars; the space constant of the interior bar was randomly scaled, elevating thresholds for the interval discrimination by a small amount (some 1.5%). This makes unlikely any strategies for computing interval width such as adding up the separation of zero-crossings in the response of a filter on a given scale, which would predict far higher thresholds than were obtained; a mechanism like the coincidence detectors of Morgan and Regan (1987) would explain the data obtained under these circumstances, since these have widely separated paired receptive fields and do not need to make intermediate measurements to assess separation. Spatial interval thresholds were also determined in the face of perturbations in the orientation of the stimuli. No clear dependence of threshold on orientation could be discovered in the experiments devised, leading to the conclusion that spatial interval judgments are rotation-invariant.

Unless MIRAGE is modified to take account of data from experiments such as the third experiment in this chapter, it cannot explain why interval discrimination should be so robust to changes in the spatial configuration of the target. One modification which was suggested by Watt (1988) is to have adaptive control over the largest filter active; if the system 'knows' that the output of low frequency filters is variable in an uninformative way, then it may be able to shut down the larger filters. This adaptive control could be a component of an attentional system.

In MIRAGE, if there are larger filters enabled, then the information contained in their resopnse cannot be selectively accessed (Morgan & Watt, 1984; Watt & Morgan, 1984). Conversely, smaller filters cannot be accessed independently of any larger filters active; (explaining why edge blur discrimination is better for moderately blurred edges than sharper edges- Watt & Morgan, 1984); but if the output of larger filters is randomly disturbed, then the system may ignore these larger filters, and collapse over a smaller frequency range of filters. This leads to the prediction that edge blur discrimination should actually *improve* at small edge blurs if some spatial perturbation which disturbs the operation of the larger filters is carried out.

Appendix: Stimuli from chapter five

Experiments 1-2: interval discrimination

Luminance profiles for the truncated Gaussian stimuli were given by:

$L(x) = L_b[1 + c.w(x)]$

L_b = background luminance,

c = contrast,

w(x) is identical to that given for stimuli from chapter 4, experiment 3; see appendix, chapter 4.

Experiments 3 and 4 had similar intensity profiles to the above, with the exception that a third bar was added for experiment 4 between the other two bars, which was itself a truncated Gaussian. In experiment 3, the inner and outer section space constants were equal; the separation of the inner zero-crossings of the bar was kept fixed throughout one experimental condition, and this was used to define the separation of the bars. Examples of intensity profiles and greyscale representations for these stimuli can be seen overleaf. The greyscale representations should be viewed from about **5 metres.**

Experiment 5: interval discrimination for stimuli of different orientations

Luminance profiles for the targets were given by:

$L(x) = L_b[1 + c.w(x)]$

where L_b = background luminance,

c = contrast,

$w(x) = g(x) + h(x)$,

where g(x) and h(x) are Gaussian functions with space constants of 1.0 min arc, and means separated by 16 min arc for the condition where the reference orientation was randomised; alternatively, the means had a median separation of 12 min arc, and were in the range 8-16 min arc, for the condition where the reference separation was randomly chosen.

Chapter 5 stimuli

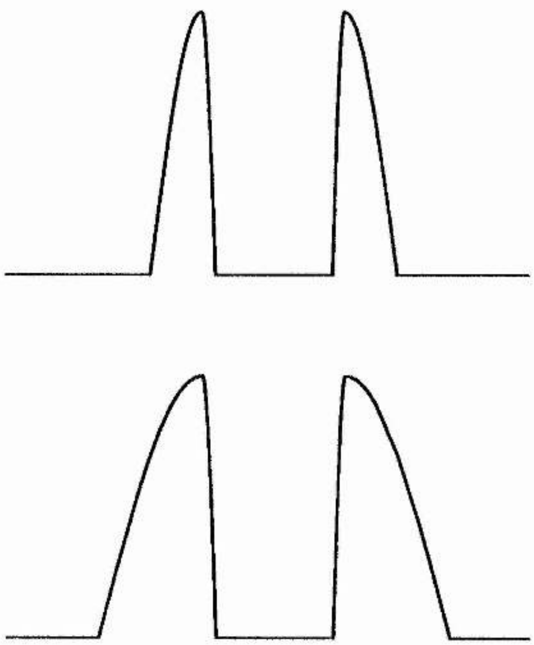

Chapter 5 stimuli

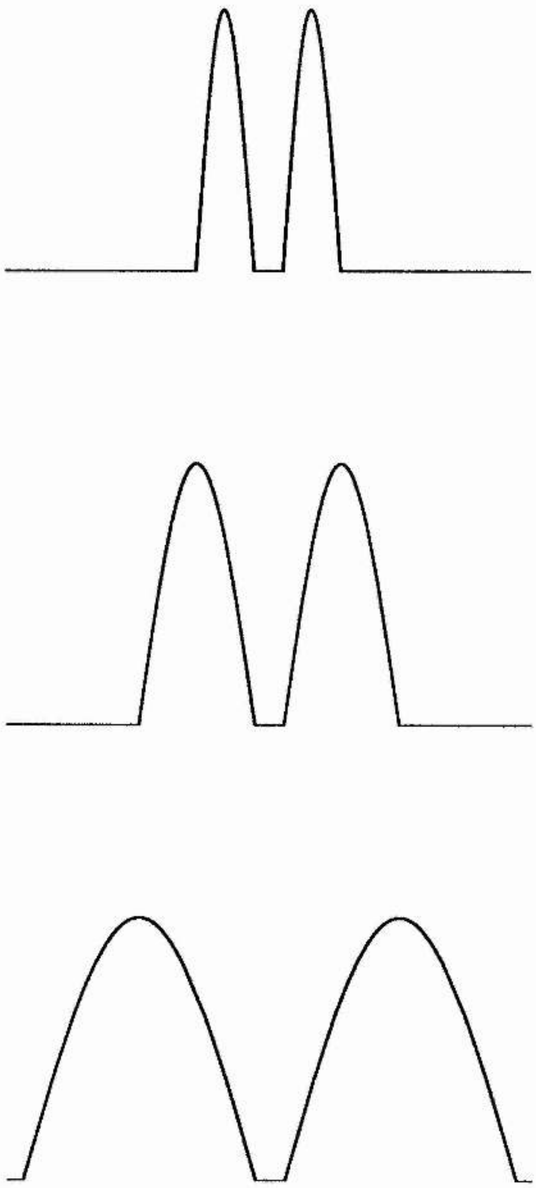

stimuli

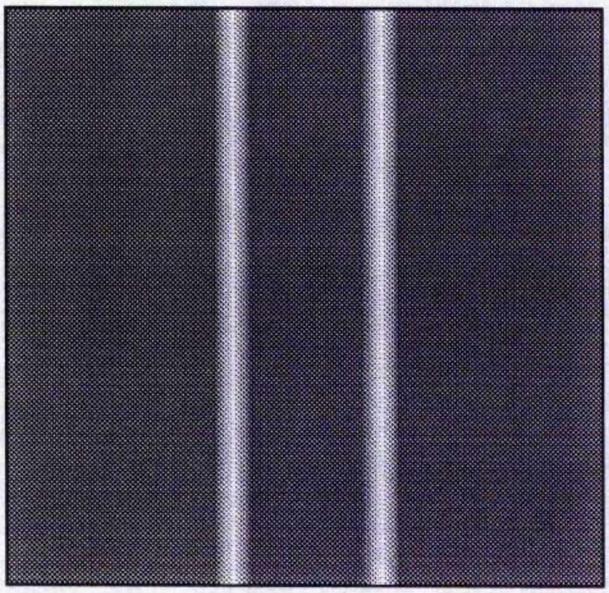

stimuli

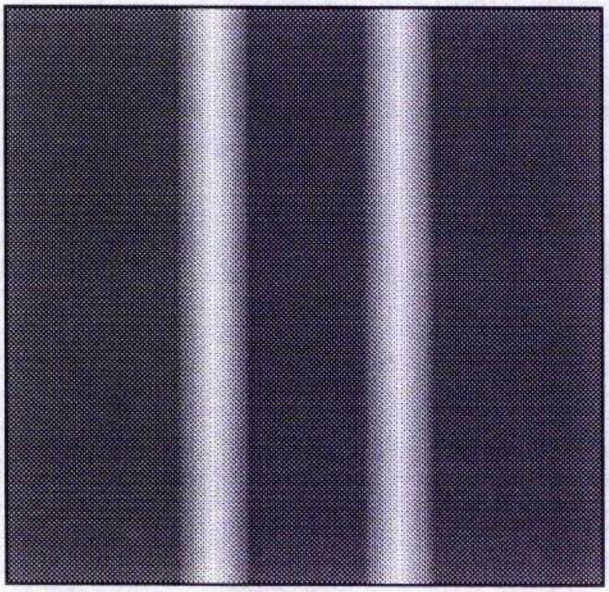

stimuli

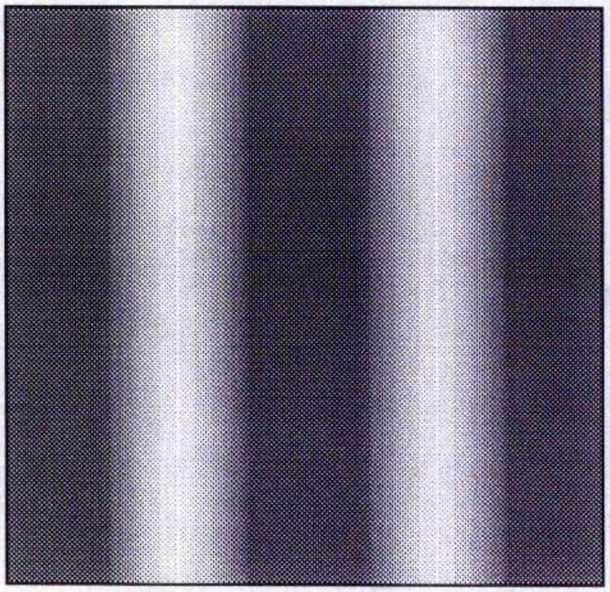

Chapter 5 stimuli

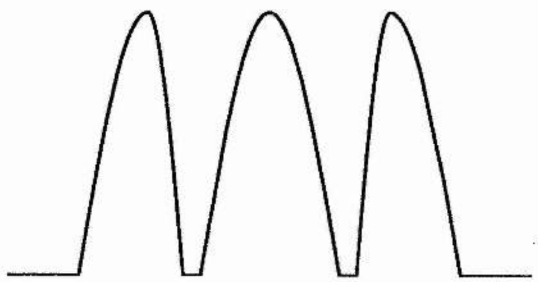

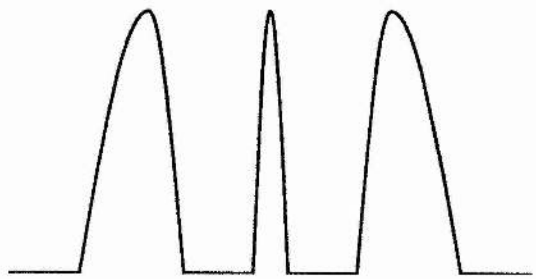

stimuli

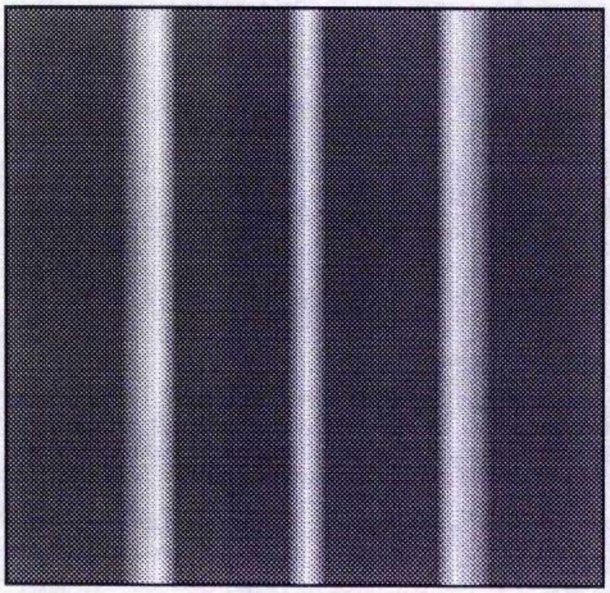

stimuli

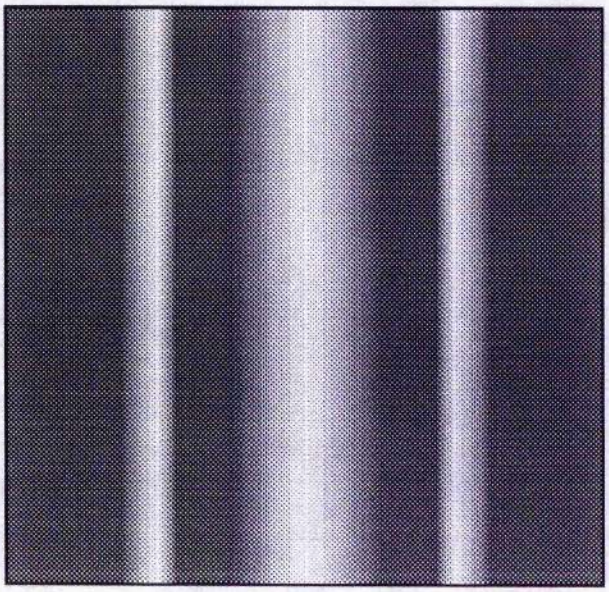

Chapter Six

Summary and conclusions

6.0: Concluding remarks

This study has explored some aspects of the relationship between the stimulus contrast profile, and the relative spatial location assigned to nearby image features. The key aspect of the experiments and modelling reported here was to reconcile the "appearance" of a visual stimulus with an account of the mechanisms behind extraction of information useful for visual tasks such as relative localisation. Introduction of a constant error term (bias) to the system's measurements of spatial location does not necessarily affect the precision with which location measurements can be made, unless the bias becomes extremely large. It was shown in chapter 3 that the sorts of measurements made in estimating spatial location with a MIRAGE algorithm introduce biases of the same sort as visual perception.

Even when adjacent features can be easily resolved, their spatial locations are distorted systematically, and this distortion is exactly consistent with the combination of relatively coarse and relatively fine information implied by the MIRAGE algorithm (Watt & Morgan, 1985). Even though the features appear as spatially disjoint, at some stage in their processing they are grouped together in such a way as to corrupt the exact locations of each feature. One such grouping mechanism is a filter large enough to blur out the spatial separation of the two features. Whilst the information used to perform the task is still available to the subject largely undegraded (since the precision on the task was unimpaired by the presence of the flanking feature in the experiments reported here), the shifts in perceived location suggest that something about the "appearance" of the target was altered by the presence of the flanking feature; the local geometry is apparently warped, and this is consistent with the distorting influence of combining spatial location estimates from different scales of filters.

It was shown that the visual system can use different strategies to localise features depending upon the task demands; if the stimulus does not vary radically in appearance across trials, then localisation is consistent with combination of information from different scales of filters. Locating the centroid of the combined response distribution of several scales of filters, and assessing spatial separation by measuring the distance between these centroids provides a good account of all the data from chapter 4, and further implies combination of information across a range of spatial scales in certain situations.

However, this strategy is not adequate when the stimulus is manipulated in such a way as to randomly perturb the location of the centroid of response of several scales of filters combined according to MIRAGE. In the presence of random variation in the orthoaxial contrast profile of the stimuli, it was clear that the visual system could selectively employ only information present on a fine scale to perform the separation judgment, consistent with the observation of Morgan and Ward (1985). One might conclude, on this basis, that the "appearance" of a stimulus reflects something about the physical properties of the stimulus across a wide range of scales. Good correspondence was found between the predictions of a simple model based on spatial primitives produced by the MIRAGE algorithm and relative location of closely spaced features subject to spatial interference. There might other explanations of the phenomenon, but the conclusion that these spatial interactions result from combination of information across spatial scale is parsimonious and consistent with previous work of Watt and Morgan.

One limitation of the study was that it was restricted to processing of foveal patterns, and did not investigate the role of eccentric presentation, which could have given further constraints upon the processes and mechanisms involved in localising intensity changes. Little attempt was made to systematically explore the relationship between contrast and the perceived width of spatial features, which has received only little attention in the literature to date; this could be interesting if it could be shown that a clear relation existed between the range of filters active and the precision of a visual judgment or perceived physical properties. To date, only speculations have been made, and the apparent insensitivity of spatial interval judgments to contrast randomisation suggest that this might not be revealing.

6.1: Main results and conclusions

The first series of experiments, described in chapter three, involved measuring biases in location for luminance features in the presence of nearby irrelevant 'flanking' features. The precision of relative localisation may be reduced with closely spaced features (Westheimer & Hauske, 1975). The experiments reported here involved measuring shifts in perceived location which were obtained without an increase in relative localisation thresholds. Shifts in perceived location were determined by measuring the shift in the location of the mean (50% point) of the subject's error response distribution.

The attraction component of spatial interference with vernier acuity and direction of displacement discrimination was found to disappear if the targets were blurred beyond a criterial blur, and perceptual attraction was observed with a displacement task at blurs of bar and retinal distances greater than a vernier acuity task. If blurred edge targets are to be localised, the effects of the polarity of contrast of the flanking feature is the opposite to the situation when a bar is to be localised; ie. an edge target is attracted to a dark flank feature, but repelled by a bright flank feature, at small target-flank separations. Spatial interference with interval width perception indicates a similar pattern of results to interference with vernier acuity for bars with bright outer flanking bars increasing the perceived width of an interval (attraction) and dark bars decreasing its perceived width (repulsion) at small target-flank distances.

The data from spatial interference experiments is well accounted for by MIRAGE. The pattern of shifts in relative location when features are closely spaced are predicted by MIRAGE. If localisation of fine lines or bars were based on the output of a single scale of filter (in other words, if the system could independently access the outputs of filters of different spatial scales), then repulsion effects would be observed at very small separations for bar targets and flanks, for example. To explain why attraction should be seen, combination of information across filters of more than one spatial scale needs to occur. The MIRAGE operation correctly explains the results of several experiments performed; it predicts the disappearance of the attraction between resolved bars in a vernier task if the bars are blurred beyond a criterial blur, and with an octave shift downwards in the peak frequencies of the filters involved, it accounts for data from an apparent motion task. It also provides a remarkably good account of the influence of flanking features on the perceived location of a target edge, in which the pattern of spatial interactions are reversed compared to those seen when the target is a bar. Modelling of the perceived shifts caused in spatial interference with relative localisation therefore provides new evidence supporting the existence of an inflexible method of combination of the responses of filters of different scales in early vision, as proposed by Watt and Morgan (1984).

A further series of interval discrimination tasks measured the perceived width of spatial intervals for bars with asymmetric luminance profiles, relative to spatial intervals delimited by symmetric bars. Perceived separation for these stimuli is not determined by features in the retinal intensity distribution, but can easily be explained by MIRAGE. A model based on only one scale of filter, taking the zero-crossings in the filter output (Toet, Smits, Nienhuis & Koenderink, 1988) could also account for these results. These experiments were not sufficiently discriminating to allow a choice to be made between competing models, although they were suggestive.

The previous series of experiments involved measuring biases in location, which were generally obtained without corresponding changes in the precision of the psychophysical judgments involved. Criticism of such experiments has been made by Watt (1990), that these measurements tell us very little that is robust about visual processing. Consequently, an attempt was made to compare the results of the previous experiments, with measurements of precision of psychophysical judgments in the face of stimulus noise, using roughly similar stimuli. Measuring precision allows us to assess both the information loss in the system as a function of external noise (which can allow us to assess the variance in the system's internal representations), and obtain measurements which are bias-free and do not depend on the vagaries of individual observers. A series of experiments were performed which involved measuring precision for a spatial interval discrimination task, with several different stimulus configurations. These experiments were described in chapter five.

Interval width discrimination thresholds were collected for a variety of spatial interval targets. The interval targets were delimited by bars composed of truncated Gaussian luminance distributions, with two distributions of (typically) different space constants being united at their luminance maxima to create a bar with an asymmetric orthoaxial contrast profile. It was found that thresholds for interval width for these stimuli are independent of the space constant of the outer section of the bars delimiting an interval (that is, the section outside the luminance maxima). In other words, interval discrimination is not affected by the degree of asymmetry of the luminance profile of the features delimiting the interval, if the inner section of the luminance profile of the features remains constant. If interval discrimination depended on localising the centroids of the intensity profile of such stimuli, then a Weber's law like relationship between the separation of the centroids and thresholds would be obtained (ie. thresholds would be proportional to the separation of the centroids of the intensity profile).

When the outer section space constant is randomly varied, ie. the degree of asymmetry of the luminance profile is varied randomly from stimulus to stimulus (but the separation of luminance maxima and inner profile remain the same) then a very small elevation of threshold was found. This was repeated for stimuli presented with an exposure duration of 20 msec, and the pattern of results obtained was virtually identical; thresholds were elevated only slightly, despite considerable variations in the width of the bars in the target. The perturbation of the space constant of the outer section randomly disturbs the output of low spatial frequency filters . Modelling of the data indicates clearly that interval width is not based on information in the outputs of low spatial frequency filters, replicating previous results (eg. Burbeck & Yap, 1990; Morgan & Ward, 1985). A MIRAGE-type approach has problems if a coarse-to-fine analysis in early vision is assumed, since no disparities in the effects of this manipulation were found between the 20 msec and 500 msec exposure durations. The representation of spatial intervals was further constrained by the results of the next experiment. Interval thresholds were collected for intervals flanked by symmetric bars, whose space constants were either fixed, or randomly varied. Interval thresholds were found to rise with a power-law function of the space constant of the bars delimiting the interval, with an exponent in the range 0.35-0.57. The data from two different separations both suggested that these functions reflected a Weber's law type of relationship between thresholds obtained and the separation of features in the neural response distribution elicited by the stimulus. For the random bar width condition, the separation of the inner edges of each bar (ie. zero-crossings in the intensity profile) remained constant, but the location of the luminance maxima was varied randomly across a range of up to 24 min arc. Several candidate features in a neural response distribution were examined, to see which could predict the appropriate pattern of increases in threshold as a function of the stimulus noise. Only features like the centroids (of a noise-truncated region) or extrema in the positive portion of the response of a very small filter (eg. a Laplacian of a Gaussian with a space constant less than 0.5 min arc) could account for the data obtained.

Interval thresholds were determined for three bar targets, the inner bar always being in the middle of the interval between the outer two bars. Randomly varying the space constant of the inner bar elevated thresholds for discrimination of the separation of the outer two bars, but the effects were small, the threshold elevation being about 1.5% (expressed as a fraction of the interval width); randomly varying both the width of the interior bar and the width of the outer section of the target bars (as in the experiment mentioned above) still allowed Weber fractions of the order of 6-8% to be obtained, compared with baselines of about 3-4%. This result was surprising, since the target's appearance from trial to trial was highly variable, yet subjects could still reliably judge the interval width between the inner sections of the outer two bars in the target. This precludes the use of certain strategies by the visual system to compute the separation of non-adjacent features; for instance, separation cannot be assessed by adding the separation of pairs of zero-crossings on a given scale within the interval between the features, since this would predict far greater loss in precision for this judgment than was obtained. The effects of perturbing the orientation of a spatial interval stimulus was also examined. It was found that interval discrimination is highly robust to perturbations in the orientation of the stimuli to be compared, but some small residual effect of orientation differences on the precision of the judgment could be found. Care was taken to avoid artefacts which might have been responsible for earlier reports of the rotation invariance of spatial period discrimination.

Features in the stimulus intensity distribution (its maxima, inflexions, threshold edges and mean) are not perceptually significant in encoding of intervals between luminance features. Modelling of the data suggests that under certain circumstances (eg. during spatial perturbation of the stimulus configuration), separation can be computed from features in the response of a very high frequency filter; MIRAGE was able to explain much of the data from the experiments without spatial perturbation. It may therefore be the case that under conditions where information in the lower frequency filters is varying in a manner unrelated to performing a task, information present in these filters can be ignored by the visual system. This could be achieved by a top-down adaptive control process, which can determine in advance the range of filters to be active, as suggested by Watt (1988). Under conditions where the outputs of large filters have noise unrelated to the task, the system may be able to switch these low spatial frequency filters out, rather than inextricably combining their output with the responses of the higher spatial frequency filters. This indicates that even the most elementary of spatial judgments might be under top-down control, which is consistent with a result of Toet, Smits, Nienhuis & Koenderink (1988) concerning localisation of bright bars; they found an interaction between the perceived location of bars and the instructions given to the subjects.

References

Anderson EE and Weymouth FW, 1923. Visual perception and the retinal mosaic. I. Retinal mean local sign: an explanation of the fineness of binocular perception of distance. *American Journal of Physiology **64**, 561-594*

Anderson SJ and Burr DC, 1987. Receptive field size of human motion detection units. *Vision Research **27**, 621-635*

Andrews DP, 1967. Perception of contour orientation in the central fovea. Part I. Short lines. *Vision Research 7, 975-997*

Andrews DP, Butcher AK and Buckley BR, 1973. Acuities for spatial arrangement in line figures: human and ideal observers compared. *Vision Research **13**, 599-620*

Badcock DR and Westheimer G, 1985a. Spatial location and hyperacuity: the centre/ surround localisation contribution function has two substrates. *Vision Research **25**, 1259-1267*

Badcock DR and Westheimer G, 1985b. Spatial location and hyperacuity: flank position within the centre and surround zones. *Spatial Vision **1**, 3-11*

Barlow HB, 1981. Critical limiting factors in the design of the eye and the visual cortex. *Proceedings of the Royal Society of London, series B, **212**, 1-34*

Blakemore CB and Campbell FW, 1969. On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology (London) **203**, 237-260*

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

Boring EG, 1942. *Sensation and perception in the history of experimental psychology*. Appleton-Century-Crofts, New York.

Bradley A and Skottun BC, 1984. The effects of large orientation and spatial frequency differences on spatial discrimination. *Vision Research **24**, 1889-1896.*

Bradley A and Skottun BC, 1987. Effects of contrast and spatial frequency on vernier acuity. *Vision Research **27**, 1817-1824*

Burbeck CA, 1987. Position and spatial frequency in large-scale localisation judgments. *Vision Research **27**, 417-427*

Burbeck CA, 1988. Large-scale relative localisation across spatial frequency channels. *Vision Research **28**, 857-859*

Burbeck DA and Regan D, 1983. Independence of orientation and size in spatial discrimination. *Journal of the Optical Society of America **73**, 1691-1694*

Burbeck CA and Yap YL, 1990. Spatial filter selection in large scale spatial-interval discrimination. *Vision Research **30**, 263-272*

Caelli T, Brettel H, Rentschler I and Hilz R, 1983. Discrimination thresholds in the two-dimensional spatial frequency domain. *Vision Research* ***23****, 129-133*

Campbell FW and Gubisch RW, 1966. Optical quality of the human eye. *Journal of Physiology* ***186****, 558-578*

Campbell FW, Nachmias J and Jukes J, 1970. Spatial frequency discrimination in human vision. *Journal of the Optical Society of America* ***60****, 555-559*

Campbell FW and Robson JG, 1968. Application of Fourier analysis to the visibility of gratings. *Journal of Physiology (Lond.)* ***197****, 551-566*

Canny JF, 1983. Finding edges and lines in images. *Artif. Intell. Lab. MIT Camb. MA, AI-TR* ***720****, June 1983*

Carlson CR and Klopfenstein RW. 1985. Spatial frequency model for hyperacuity. *Journal of the Optical Society of America* ***A2****, 1747-1751*

Corwin CR, Moskowitz-Cook A and Green MA, 1977. The oblique effect in a vernier acuity situation. *Perception and Psychophysics* ***21****, 445-449*

Daugman JG, 1980. Two dimensional spectral analysis of cortical receptive field profiles. *Vision Research* ***20****, 847-856*

Daugman JG, 1985. Uncertainty resolution in space, spatial frequency and orientation optimised by two-dimensional cortical filters. *Journal of the Optical Society of America* ***A2****, 1160-1169*

De Valois RL, Albrecht DG and Thorell LG, 1982. Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research* ***22****, 531-544*

Finney DJ, 1971. *Probit Analysis*, 3rd edition, Cambridge University Press

Gabor D, 1946. Theory of communication. *Journal of the Institute of Electrical Engineers (London)* ***93****, 429-457*

Ganz L, 1964. Lateral inhibition and the location of visual contours: an analysis of visual after-effects. *Vision Research* ***4****, 465-481*

Ganz L, 1966. Mechanism of the figural after-effects. *Psychological Review* ***73****, 128-150*

Gelb DJ and Wilson HR, 1983a. Shifts in perceived size as a function of contrast and temporal modulation. *Vision Research* ***23****, 71-82*

Gelb DJ and Wilson HR, 1983b. Shifts in perceived size due to masking. *Vision Research* ***23****, 589-597*

Georgeson MA, 1980. Spatial frequency analysis in early visual processing. *Philosophical Transactions of the Royal Society, series B,* ***290****, 11-22*

Graham N and Nachmias J, 1971. Detection of grating patterns containing two spatial frequencies: a comparison of single channel and multiple channel models. *Vision Research* ***11****, 251-259*

Graham N, Robson JG and Nachmias J, 1978. Grating summation in fovea and periphery. *Vision Research* ***18****, 815-825*

Green DG, 1970. Regional variations in the visual acuity for interference fringes. *Journal of Physiology (London)* ***207****, 351-356*

Heeley DW, 1987. Spatial frequency discrimination for sinewave gratings with random, bandpass frequency modulation: evidence for averaging in spatial acuity. *Spatial Vision* ***2****, 317-335*

Heeley DW and Timney B, 1987. Meridional anisotropies of orientation discrimination for sinewave gratings. *Vision Research* ***27****, 337-314*

Hess RF, Pointer JS and Watt RJ, 1989. How are spatial filters used in fovea and parafovea? *Journal of the Optical Society of America* ***A6****, 329-339*

Hess RF, Pointer JS and Watt RJ, 1990. Regional distribution of the mechanisms that underlie spatial localisation. *Vision Research* ***30****, 1021-1031*

Hines M, 1976. Line spread function variation near the fovea. *Vision Research* ***16****, 567-572*

Hirsch J and Hylton R, 1982. Limits of spatial frequency discrimination as evidence of neural interpolation. *Journal of the Optical Society of America* ***72****, 1367-1374*

Hubel DH and Wiesel TN, 1962. Receptive fields, binocular interaction and functional architecture in the cat's striate cortex. *Journal of Physiology* ***160****, 106-154*

Hubel DH and Wiesel TN, 1968. Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)* ***195****, 215-243*

Hubel DH and Wiesel TN, 1974a. Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology* ***158****, 267-294*

Hubel DH and Wiesel TN, 1974b. Uniformity of striate cortex: a parallel relationship between field size, scatter and magnification factor. *Journal of Comparative Neurology* ***158****, 295-305*

Jones JP and Palmer LA, 1987. The two-dimensional spatial structure of receptive fields in cat striate cortex. *Journal of Neurophysiology* ***58****, 1187-1211*

Klein S, Stromeyer CF and Ganz L, 1974. The simultaneous spatial frequency shift: a dissociation between the detection and perception of gratings. *Vision Research* ***14****, 1421-1432*

Kohler W and Wallach H, 1944. Figural after-effects, an investigation of visual processes. *Proceedings of the American Philosophical Society* ***88****, 269-357*

Krose BJA, 1987. Local structure analysers as determinants of preattentive pattern discrimination. *Biological Cybernetics* ***55****, 289-298*

Krauskopf J and Campbell FW. The effects of contrast and target width on vernier acuity. Unpublished manuscript

Kulikowski JJ, Marcelja S and Bishop PO, 1982. Theory of spatial position and spatial frequency relations in the receptive fields of simple cells in the visual cortex. *Biological Cybernetics* ***43****, 187-198*

Laming D and Marsh D, 1986. A QUEST procedure for general use: application to the measurement of vibrotactile thresholds. Original manuscript

Levi DM, Klein SA and Aitsebaomo AP, 1985. Vernier acuity, crowding and cortical magnification. *Vision Research* ***25****, 963-977*

Levi DM and Westheimer G, 1987. Spatial interval discrimination in the human fovea: what delimits the interval? *Journal of the Optical Society of America* ***A4****, 1304-1313*

Limb JO and Rubinstein CB, 1977. A model of threshold vision incorporating inhomogeneity of the visual field. *Vision Research* ***17****, 571-584*

Ludvigh E, 1953. Direction sense of the eye. *American Journal of Opthalmology* ***36****, 139-143*

McKee SP and Westheimer G, 1978. Improvement in vernier acuity with practice. *Perception and Psychophysics* ***24****, 258-262*

Madigan R and Williams D, 1987. Maximum likelihood psychometric procedures in two alternative forced choice: evaluation and recommendations. *Perception and Psychophysics* ***42****, 240-249*

Marr D, 1976. Early processing of visual information. *Philosophical Transactions of the Royal Society, series B,* ***275****, 483-524*

Marr D, 1982. *Vision.* WH Freeman and Co. San Francisco

Marr D and Hildreth E, 1980. A theory of edge detection. *Proceedings of the Royal Society of London, series B,* ***207****, 187-217*

Marr D, Poggio T and Hildreth E, 1980. Smallest channel in early human vision. *Journal of the Optical Society of America* ***70****, 868-870*

Marr D, Ullman S and Poggio T, 1979. Bandpass channels, zero-crossings and early visual information processing. *Journal of the Optical Society of America* ***69****, 914-916*

Mather G, 1987. The dependence of edge displacement thresholds on edge blur, contrast and displacement distance. *Vision Research* ***27****, 1631-1637*

Mather G and Morgan MJ, 1986. Irradiation: implications for theories of edge localisation. *Vision Research* ***26****, 1007-1015*

Morgan MJ, Mather G, Moulden B and Watt RJ, 1984. Intensity-response non-linearities and the theory of edge localisation. *Vision Research* ***24****, 713-719*

Morgan MJ and Regan DM, 1987. Opponent model for line interval discrimination: interval and vernier performance compared. *Vision Research* ***27****, 107-118*

Morgan MJ and Ward RM, 1985. Spatial and spatial frequency primitives in spatial interval discrimination. *Journal of the Optical Society of America* ***A2****, 1205-1210*

Morgan MJ and Watt RJ, 1982. Mechanisms of interpolation in human spatial vision. *Nature (London)* ***299****, 553-555*

Morgan MJ and Watt RJ, 1984. Spatial frequency interference effects and interpolation in vernier acuity. *Vision Research* ***24***, *1911-1919*

Nakayama K and Silverman GH, 1985. Detection and discrimination of sinusoidal grating displacements. *Journal of the Optical Society of America* ***A2***, *267-274*

Nielsen KRK and Wandell BA, 1988. Discrete analysis of spatial sensitivity models. *Journal of the Optical Society of America* ***A5***, *743-755*

Olzak LA and Thomas JP, 1987. Seeing Spatial Patterns. Chapter 7, *Handbook of Human Perception and Performance.*

Orban GA, Vandenbussche E and Vogels R, 1984. Meridional variations and other properties suggesting that acuity and orientation discrimination rely on different neuronal mechanisms. *Opthalmological and Physiological Optics* ***4***, *89-93*

Paradiso MA, 1988. A theory for the use of visual orientation information which exploits the columnar structure of striate cortex. *Biological Cybernetics* ***58***, *35-49*

Parker AH and Hawken MJ, 1985. Capabilities of monkey cortical cells in spatial resolution tasks. *Journal of the Optical Society of America* ***A2***, *1102-1114*

Parker AH and Hawken MJ, 1988. Two-dimensional spatial structure of receptive fields in monkey striate cortex. *Journal of the Optical Society of America* ***A5***, *598-605*

Poggio GF, Motter BC, Squatrito S and Trotter Y, 1985. Responses of neurones in visual cortex (V1 and V2) of the alert macaque to dynamic random dot stereograms. *Vision Research* ***25***, *397-406*

Quick RFA, 1974. A vector magnitude model of contrast detection. *Kybernetik,* ***16***, *65-67*

Ratliff F, 1984. Why Mach bands are not seen at the edges of a step. *Vision Research* ***24***, *163-165*

Regan D, 1985. Storage of spatial frequency information and spatial frequency discrimination. *Journal of the Optical Society of America* ***A2***, *619-621*

Regan D and Beverley KI, 1985. Postadaptation orientation discrimination. *Journal of the Optical Society of America* ***A2***, *147-155*

Sachs MB, Nachmias J and Robson JG, 1971. Spatial frequency channels in human vision. *Journal of the Optical Society of America* ***61***, *1176-1186*

Skottun BC, Bradley A, Sclar G, Ohzawa I and Freeman RD, 1987. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behaviour. *Journal of Neurophysiology* ***57***, *773-786*

Stigmar G, 1971. Blurred visual stimuli. II. The effect of blurred visual stimuli on vernier and stereo acuity. *Acta Opthalmologica* ***49***, *364-379*

Thomas JP, Gille J and Barker RA, 1982. Simultaneous detection and identification: theory and data. *Journal of the Optical Society of America* ***72***, *1642-1651*

Toet A and Koenderink JJ, 1988. Differential spatial displacement discrimination thresholds for Gabor patches. *Vision Research* ***28***, *133-143*

Toet A, Smits CS, Nienhuis B and Koenderink JJ, 1988. The visual assessment of the spatial location of a bright bar. *Vision Research* ***28**, 721-737*

Toet A, Van Eekhout MP, Simons HLJT and Koenderink JJ, 1987. Scale invariant features of differential spatial displacement discrimination. *Vision Research* ***27**, 441-451*

Watt RJ, 1984. Towards a general theory of the visual acuities for shape and spatial arrangement. *Vision Research* ***24**, 1377-1386*

Watt RJ, 1987. Scanning from coarse to fine spatial scales in the human visual system after the onset of a stimulus. *Journal of the Optical Society of America* ***A4**, 2006-2021*

Watt RJ, 1988. *Visual Processing*. Lawrence Erlbaum, Hove.

Watt RJ, 1990. *Understanding Vision*. Forthcoming.

Watt RJ and Andrews DP, 1981. APE: adaptive probit estimation of psychometric functions. *Current Psychological Review* ***1**, 205-215*

Wat RJ and Andrews DP, 1982. Contour curvature analysis: hyperacuities in the discrimination of detailed shape. *Vision Research* ***22**, 449-460*

Watt RJ and Morgan MJ, 1983a. Mechanisms response for the assessment of visual location: theory and evidence. *Vision Research* ***23**, 97-109*

Watt RJ and Morgan MJ, 1983b. The recognition and representation of edge blur: evidence for spatial primitives in human vision. *Vision Research* ***23**, 1465-1477*

Watt RJ and Morgan MJ, 1984. Spatial filters and the localisation of luminance changes in human vision. *Vision Research* ***24**, 1387-1397*

Watt RJ and Morgan MJ, 1985. A theory of the primitive spatial code in human vision. *Vision Research* ***25**, 1661-1674*

Watt RJ, Morgan MJ and Ward RM, 1983a. The use of different cues in vernier acuity. *Vision Research* ***23**, 991-995*

Watt RJ, Morgan MJ and Ward RM, 1983b. Stimulus features that determine the visual location of a bright bar. *Invest. Ophtalmol. Vis. Sci.* ***24**, 66-71*

Watson AB, 1982. Summation of grating patches indicates many types of detectors in one retinal location. *Vision Research* ***22**, 17-25*

Watson AB and Pelli DG, 1983. QUEST: A Bayesian adaptive psychometric method. *Perception and Psychophysics* ***33**, 113-120*

Watson AB and Robson JG, 1981. Discrimination at threshold: labelled detectors in human vision. *Vision Research* ***21**, 1115-1122*

Westheimer G, 1984. Line separation discrimination in the human fovea: smooth or segmented? *Journal of the Optical Society of America* ***A1**, 683-684*

Westheimer G and Hauske G 1975. Temporal and spatial interference with vernier acuity. *Vision Research* ***15**, 1136-1141*

Westheimer G and McKee SP, 1977a. Spatial configurations for visual hyperacuity. *Vision Research* ***17****, 941-947*

Westheimer G and McKee SP, 1977b. Integration regions for visual hyperacuity. *Vision Research* ***17****, 89-93*

Westheimer G, Shimamura K and McKee SP, 1976. Interference with line orientation sensitivity. *Journal of the Optical Society of America* ***66****, 332-338.*

Wilson HR, 1978. Quantitative characterisation of two types of line spread function near the fovea. *Vision Research* ***18****, 971-981*

Wilson HR, 1986. Responses of spatial mechanisms can explain hyperacuity. *Vision Research* ***26****, 453-469*

Wilson HR and Bergen JR, 1979. A four mechanism model for threshold spatial vision. *Vision Research* ***19****, 19-32*

Wilson HR and Gelb DJ, 1984. Modified line-element theory for spatial frequency and width discrimination. *Journal of the Optical Society of America* ***A1****, 124-131*

Wilson HR, MacFarlane DK and Phillips GC, 1983. Spatial frequency tuning of orientation selective units estimated by oblique masking. *Vision Research* ***23****, 873-882*