INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

UMI[®]

NOTE TO USERS

This reproduction is the best copy available.

UMI

**Disparity Contingent High Spatial Frequency Constraints
on the Upper Velocity Limit of Stereopsis**

by

Stan S. Lee, B. A., M. A.

**A thesis submitted to
the Faculty of Graduate Studies and Research
in partial fulfilment of
the requirements for the degree of
Doctor of Philosophy**

Department of Psychology

**Carleton University
Ottawa, Ontario
May, 2001**

© copyright

2001, Stan S. Lee

**National Library
of Canada**

**Acquisitions and
Bibliographic Services**

**395 Wellington Street
Ottawa ON K1A 0N4
Canada**

**Bibliothèque nationale
du Canada**

**Acquisitions et
services bibliographiques**

**395, rue Wellington
Ottawa ON K1A 0N4
Canada**

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-60961-8

Canada

The undersigned recommend to the
Faculty of Graduate Studies and Research
Acceptance of the thesis:

"Disparity Contingent High Spatial Frequency Constraints on the Upper Velocity Limit
of Stereopsis"

submitted by

Stan Lee

in Partial fulfillment of the requirements for

the degree of Doctor of Philosophy

Chair

External Examiner

Thesis Supervisor

Carleton University
April, 2001

ABSTRACT

Three experiments were conducted to examine the nature of spatial frequency interactions in the processing of binocular disparities in lateral motion. A recent report by Morgan and Castet (1995) suggested that the upper temporal frequency limit of stereopsis was 30 to 40 Hz regardless of spatial frequency. However, the observation of a constant upper temporal frequency limit may have resulted due to the restricted range of spatial frequencies used in their experiment (0.04 to 0.32 cycle/degree). The first experiment demonstrated that the upper temporal frequency limit of stereopsis is indeed contingent on spatial frequency where higher spatial frequencies exhibited progressively poorer temporal resolution. These results are in agreement with the notion that disparity is processed in channels which are tuned to spatial and temporal frequencies. In an effort to identify interactions across spatial frequencies in stereopsis, a second experiment was conducted where spatial frequencies processed with similar and different temporal resolutions were combined in compound gratings. Upper velocity limits for compound gratings comprising spatial frequencies processed with differing temporal resolutions yielded a reduction in the upper velocity limit consistent with the presence of high spatial frequency constraints. In a third experiment I demonstrate that these constraints are more likely to be observed at small disparities and are largely absent at larger disparities. This suggests that the magnitude of disparity plays an important role in determining how motion signals are bound.

ACKNOWLEDGEMENTS

Sincere thanks to both of my advisors, Lew Stelmach and Chris Herdman for their valuable help and guidance as well as patience. I would also like to thank the members of my examining committee: Dr. Logan, Dr. Tansley, Dr. Bischof, and Dr. Brook. But most of all I would like to thank Millie who spent many a night by me during my academic years, Zak and Chloé to whom I owe hours of well needed distraction. Finally, to my wife Nathalie to whom I owe so much.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vi
INTRODUCTION	1
Stereo Acuity for Objects in Motion	4
The Morgan and Castet Study	8
Low to High Spatial Frequency Constraints on Stereopsis	14
SUMMARY	20
High to Low Spatial Frequency Constraints on Stereopsis	22
Stereopsis for static Images	22
Disparity and Motion	29
SUMMARY	36
EXPERIMENT I	39
METHOD	40
Observers	40
Apparatus	40
Stimuli and Procedure	42
RESULTS	47
Stereo Thresholds for Static Displays	47
Stereo Performance for Drifting Gratings	49
DISCUSSION	58
EXPERIMENT II A	60
METHOD	62
Stimuli and Procedure	62
RESULTS	62
Static stereo thresholds	62
Dynamic stereo thresholds (adjacent compound gratings)	64
Dynamic stereo thresholds (non-adjacent compound gratings)	69
EXPERIMENT II B	70
RESULTS	71
DISCUSSION	75
Signal to Noise Ratio Due to Display Aliasing	79
Size-Disparity Correlation	81

EXPERIMENT III	88
METHOD	88
Stimuli and Procedure	88
RESULTS	89
Static Disparity Thresholds	89
Simple Grating	89
Adjacent Compound Gratings	93
DISCUSSION	99
 GENERAL DISCUSSION	 101
Ecological Perspective	103
Suggested Experiments	109
Inter Ocular Contrast Differences	110
Disparity Averaging	111
Motion Transparency	111
One Final Point : High Spatial Frequency Constraints in Motion Revisited	 113
 References	 119
 Appendix 1	 126

LIST OF FIGURES

Figure 1 :	Intensity values (cd/m^2) as a function of Look Up Table (LUT) index	41
Figure 2 :	Sample experimental display depicting simple grating used In Experiment I	44
Figure 3 :	Static disparity thresholds (min) as a function of spatial frequency. Error bars represent two standard errors	48
Figures 4 a and b :	Relative disparity thresholds (simple gratings) as a function of velocity	50
Figures 4 c and d :	Relative disparity thresholds (simple gratings) as a function of velocity	51
Figure 5 :	Upper velocity limit (hollow symbols) and upper temporal frequency limit (solid symbols) as a function of spatial frequency (data averaged across four observers). Error bars represent one standard error	54
Figures 6 a and b :	Relative disparity thresholds (simple gratings) as a function of temporal frequency	56
Figure 6 c and d :	Relative disparity threshold (simple gratings) as a function of temporal frequency	57
Figure 7 :	Static disparity thresholds (min) for simple (hollow symbols) and compound (solid symbols) gratings as a function of spatial frequency. Error bars represent two standard errors	63
Figures 8 a and b :	Relative disparity thresholds for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity	65
Figures 8 c and d :	Relative disparity thresholds for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity	66
Figure 9 :	Average (four observers) upper velocity limit for simple (black bars) and adjacent compound (white bars) gratings. Error bars represent two standard error	68

Figures 10 a and b :	Relative disparity thresholds for simple (solid symbols) and compound peak subtracted (hollow symbols) gratings as a function of velocity	72
Figures 10 c and d :	Relative disparity thresholds for simple (solid symbols) and compound peak subtracted (hollow symbols) gratings as a function of velocity	73
Figure 11 :	Average upper velocity limit (four observers) for simple (black bars), peak added compound (white bars), and peak subtracted compound (grey bars) gratings. Error bars represent two standard errors	74
Figure 12 :	Static disparity thresholds for 75% and 100% correct responses Error bars represent two standard errors	90
Figures 13 a and b :	P correct (hollow symbols) (Exp III) and relative disparity thresholds (solid symbols) (Exp I) as a function of velocity	91
Figures 13 c and d :	P correct (hollow symbols) (Exp III) and relative disparity thresholds (solid symbols) (Exp I) as a function of velocity	92
Figures 14 a and b :	P correct as a function of temporal frequency	94
Figures 14 c and d :	P correct as a function of temporal frequency	95
Figures 15 a and b:	P correct (primary Y-axis) and relative disparity threshold (secondary Y-axis) for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity.	96
Figures 15 c and d:	P correct (primary Y-axis) and relative disparity threshold (secondary Y-axis) for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity.	97
Figure 16 :	Observers often fixate an object as they navigate in the visual environment	108

INTRODUCTION

The eyes, on average, are laterally separated by 6.5 cm and therefore see the visual world from different angular perspectives. As a result, the retinal projections of our visual environment are also slightly displaced on the horizontal extent. Such retinal position displacements are called binocular disparities. In order for observers to view an object foveally, observers converge their eyes onto its location, effectively bringing together the disparate retinal projections. By converging the eyes in such a way, the visual system allows all objects falling at that same distance to project onto corresponding retinal points (the collection of these points form a fixation plane called the horopter). Hence, by definition, objects falling onto the horopter will have a disparity of zero, whereas objects falling in front of, or behind, the horopter will have a disparity other than zero.

The ability of the visual system to compute the absolute distance of an object, with respect to the observer, does not mean that it is unaware of the distance of objects not falling onto the fixation plane. In fact, because the disparity between stimulated retinal points is geometrically related to the distance of objects from the horopter, the visual system is able to derive the relative depth of objects with respect to fixation. Furthermore, the visual system can also derive whether the object is lying in front of, or behind, the horopter and this even in the absence of monocular cues. This is achieved by considering whether or not the disparity between stimulated

retinal points is cross or uncrossed with objects having crossed disparities appearing in front, and objects having uncrossed disparities appearing behind the horopter.

The idea that binocular disparities contribute critically to the perception of depth derives from the invention of the stereoscope by Wheatstone in the 19th century. It is Wheatstone who first demonstrated that the visual system uses horizontal disparities to establish the relative depth of objects with respect to the fixation point. Since Wheatstone's demonstration, stereoscopic perception has been the subject of extensive studies in the fields of psychophysics (Dev, 1975; Julesz, 1960, 1971; Marr & Poggio, 1976, 1979; Mayhew & Frisby, 1981) and neurophysiology (Pettigrew, 1965 : as cited in Barlow, Blakemore, & Pettigrew (1967); Barlow, Blakemore, & Pettigrew, 1967; Hubel & Wiesel, 1970; Poggio, 1984; Hubel & Livingston, 1987). Perhaps the most seminal psychophysical work on disparity and stereopsis was conducted by Julesz in 1960 when he introduced random-dot stereograms. Random-dot stereograms consist of two images (one presented to the right and the other to the left eye) initially composed of an identical pattern of random texture. Within one of the images, however, a subset of the texture depicting a form (e.g., a square) is shifted to the left or right. Because each image in random-dot stereograms is composed of random texture, no form can be seen within each image by itself. Only when the images are viewed stereoscopically the portion of dots, which is shifted, can be seen as appearing either in front of, or behind, the remaining texture depending on whether

the shift is to the left or to the right. This is important because it provides conclusive evidence that binocular disparities are used by the visual system to infer depth even in the absence of a recognizable monocular form. The impact that Julesz's demonstration had on stereopsis research can not be overstated as the following decades produced a substantial amount of interest in the identification of the neural substrate and mechanisms that are responsible for encoding binocular disparities. For example, although no formal mechanism was proposed in the original description of the binocular function of the visual cortex, Hubel and Wiesel (1962) had already speculated that cortical cells may be involved in encoding depth information. It is difficult to imagine that such a speculation would have been possible without the important contribution of Julesz (1960).

Although the potential existence of cortical neurons responsible for the encoding of binocular disparities was first established by Hubel and Wiesel (1962), the first direct physiological evidence of disparity encoding in the brain was only obtained in the mid-1960s when Pettigrew (1965) (as cited in Barlow, Blakemore, & Pettigrew, 1967) conducted electrophysiological studies of binocular neurons in the striate cortex of the cat. Results from his experiments demonstrated the existence of disparity selective cells in the primary visual area. In a later study, Barlow, et al., (1967) extended these results by presenting, for the first time, the notion that there are binocular neurons in the primary visual cortex which respond selectively to stimuli

positioned at various distances from the eyes (i.e., different retinal disparities). These results (Barlow et al., 1967) ultimately lead to the widely accepted notion that binocular cells respond to different depths in space relative to the fixation point (i.e., different binocular disparities). Barlow et al., (1967) also suggested that this coding was specialized for horizontal disparities. This suggestion is based on the finding that there are neurons with receptive fields that possess different positional offsets from the corresponding retinal point, and that the distribution of preferred disparities was larger for the horizontal dimension than for the vertical.

Given the psychophysical data gathered in behavioural experiments and the neurophysiological evidence for the existence of a neural substrate for the encoding of depth information through binocular disparities, it is not surprising that most contemporary theories of stereopsis rely heavily on the computation of disparities to explain depth perception (Dev, 1975; Marr & Poggio, 1976, 1979; Mayhew & Frisby, 1981; Anzai, Ozhawa, & Freeman, 1999; DeAngelis, Ohzawa, & Freeman, 1995). Although these models differ with regard to how the computation of binocular disparities is achieved, disparity has and continues to drive the crux of most models of stereopsis.

Stereo Acuity for Objects in Motion

In order to assess the limits of human ability to perceive relative depth, experimenters have often relied on the smallest disparity at which observers can still

reliably report depth between two stimuli (one at zero disparity and the other at a non-zero disparity). The general agreement resulting from these studies is that, under the best conditions, stereo acuity is somewhere in the range of 2 to 6 arc sec of disparity, which is similar to that observed with a standard Vernier acuity task (Tyler, 1973; Howard, 1919; Anderson & Weymouth, 1923, Woodburn, 1934). As noted by Morgan and Castet (1995), however, studies that have examined the lower disparity limits of stereopsis have been mainly restricted to static images. This is a problem because our visual environment is not strictly limited to static images. In addition, observers often navigate in their natural setting causing even static objects to produce dynamic retinal projections.

A number of studies have examined the ability of observers to accurately perceive depth for objects in motion (Richards, 1951; Tyler, 1971; Regan & Beverly, 1973; Tyler, Schor, & Colletta, 1992; Lankheet & Lennie, 1995). Overall, the results indicate that the binocular visual system is rather sluggish in that disparity information was lost at temporal frequencies above 4 Hz to 6 Hz. It is somewhat surprising that the visual system would be so sluggish in determining dynamically the depth of objects in motion. Indeed, these results suggest that the visual system employs binocular disparities mainly for the encoding of objects moving at low velocities, and that it relies solely or largely on monocular depth cues for objects moving beyond 4 Hz to 6 Hz. This is somewhat unlikely given the rather large proportion of binocular

cells that demonstrate disparity selectivity (an estimated 66% of cells in area MT) (Maunsell and Van Essen, 1983).

As pointed out by Morgan and Castet (1995), these studies may not provide a complete measure of the temporal properties of the visual system in encoding disparities in dynamic displays. This is due to the fact that these studies have examined the spatiotemporal aspects of stereoscopic vision using experimental procedures which temporally modulate the disparity signal. The studies' findings may thus be limited to objects moving towards, or away from, the observer and may only represent an indication of the limitations of the visual system to assess a dynamic gradient of disparities. In the visual environment, however, the motion of objects is not necessarily restricted to movement towards or away from observers and may include a lateral component as well.

Lateral displacements of objects on the retinae pose a different problem with respect to dynamic disparity computations. In contrast to motion in depth, the retinal disparities of laterally moving objects can remain relatively the same regardless of the position of the object. Thus, instead of stimulating a variety of disparity selective neurons, as occurs with motion in depth, the detection of disparity in lateral motion relies mainly on the temporal resolution of disparity specific neurons. Consider, for example, a motor vehicle moving on a street perpendicular to you while you are looking at a sign directly across from this same street. For the sake of simplicity, let

us further assume that the vehicle is moving at a constant velocity. Because the vehicle is located away from fixation (i.e., at a depth different than that of the sign upon which your eyes are converging), it will mainly stimulate non-corresponding retinal points. However, the fact that it is moving implies that it will eventually stimulate corresponding retinal points, albeit with a given temporal delay. In order to accurately assess the relative depth of the vehicle, the visual system must be able to decipher the temporal delay between the stimulation of corresponding retinal points. As another example, a bar moving laterally at a fixed velocity of 100° per second at a constant disparity of 0.1° , will stimulate corresponding retinal points with a temporal delay of 1 msec. If the visual system were unable to decipher a signal arising from a 1 msec phase difference, the bar would be perceived as having stimulated corresponding retinal points at the same time, yielding a zero disparity signal. From this, it becomes obvious that, for laterally moving objects, the sensitivity of the visual system to disparity signals is limited by the temporal resolution of the receptive fields of disparity detectors. In other words, the extent to which a disparity selective neuron is able to assess the relative depth of a laterally moving object is limited by its temporal resolving power (Morgan & Castet, 1995). Consequently, there should be a maximum velocity at which stereopsis will function efficiently. Despite this observation, little has been done to assess the upper velocity limit of stereopsis. In fact, an extensive review of the literature revealed only the study of Morgan and Castet.

The Morgan and Castet Study

In their experiment, Morgan and Castet asked observers to view vertically oriented gratings of 0.04, 0.08, 0.16, and 0.32 cycle/degree which were drifted horizontally at various velocities. The benefit in using such spatial frequency gratings is that they allow an independent manipulation of velocity (expressed in degrees per second) and temporal frequency (expressed as cycles per second or Hz), where temporal frequency of a grating is defined as the number of cycles passing a given point in space per second, and velocity is defined as the temporal frequency divided by the spatial frequency. Hence, for a given velocity, temporal frequency could be manipulated by varying the spatial frequency of the grating (a higher spatial frequency grating yielding a higher temporal frequency, and a lower spatial frequency grating a lower temporal frequency).

The displays used by Morgan and Castet consisted in a pair of identical gratings that were presented on the left and right half of a computer monitor. Stereo disparities were introduced by shifting the position of one of the gratings (i.e., advancing or retarding the phase of the right eye grating relative to the left). This caused the fused grating to appear either in front of, or behind, a fixation plane consisting of a static random-dot pattern presented at zero disparity. The task of the observer was to report whether the grating had appeared in front of, or behind, the surrounding static random-dot pattern. Using static versions of their displays, Morgan

and Castet determined the disparity threshold for each observer at each spatial frequency. The disparity thresholds were defined as the amount of phase shift (disparity) needed to achieve an 84% correct criteria. Using an adaptive psychometric procedure, Morgan and Castet then assessed the observers' ability to maintain the disparity threshold at increasing velocities.

Morgan and Castet's results demonstrated that, contrary to temporally modulated disparities, stereo acuity for laterally moving objects could be maintained for very high velocities (i.e., 640° per second in the case of the lowest spatial frequency) which decreased proportionately with increasing spatial frequency. However, as velocity is proportional to the spatial frequency, Morgan and Castet argued that the true limit of stereopsis was not defined by velocity at all but by the temporal frequency of the display¹. Indeed, when the data were plotted against temporal frequency instead of velocity, stereo acuity broke down at approximately 30 to 40 Hz for all spatial frequencies. Based on these results, Morgan and Castet argued

1

The proportional relationship between the upper velocity limit and spatial frequency is due to the geometrical constraints imposed by differing spatial frequencies. With periodic stimuli such as spatial frequency gratings, the maximum displacement over which motion can be unambiguously perceived is half the spatial period (i.e. 180°). As spatial frequency increases however, a 180° displacement results in visual angle displacements which proportionately decrease as spatial frequency increases. Thus, although observers may be able to resolve a 180° displacement at all spatial frequencies, the resulting velocity decreases proportionately with increasing spatial frequency. Upper velocity limit data can be converted as upper temporal frequency limits by multiplying the velocity with the appropriate spatial frequency. This results in a measure of motion which takes into account the spatial frequency of the stimulus. Thus, equivalent temporal resolutions across spatial frequencies are reflected as a proportional reductions in the upper velocity limit or alternatively, a constant upper temporal frequency limit.

that the temporal resolution of the visual system for processing disparities in lateral motion was approximately 30 to 40 Hz regardless of spatial frequency.

There are physiological considerations as well as psychophysical data suggesting that the decline of stereoscopic performance between 30 and 40 Hz as observed by Morgan and Castet (1995) may only represent a partial picture of the upper limit of the visual system's ability to compute depth in lateral motion. This is owing to the restricted range of spatial frequencies used by Morgan and Castet (1995). For example, there is strong evidence that the visual system processes visual information through channels that are tuned to spatial frequency (Campbell & Robson, 1968; Blakemore & Campbell, 1969; Graham, Robson, & Nachmias, 1978; Shoham, Hubener, Schulze, Grivald, & Bonhoeffer, 1997; Waugh & Hess, 1994). These channels differ markedly in their temporal responses with channels tuned to high spatial frequencies preferring low temporal frequencies, and those tuned to low spatial frequencies preferring high temporal frequencies. Channels tuned to low spatial frequencies respond best to signals between 5 and 10 Hz and are capable of resolving signals up to 30 to 40 Hz (Robson, 1966).

Recent results by Shoham et al. (1997) suggest that there are differences in the temporal frequency preference of channels tuned to spatial frequencies ranging between 0.1 and 1 cycle/degree. One interesting aspect of their results is that the temporal frequency preference of neurons tuned to spatial frequencies below 0.6

cycle/degree were so similar that they argued for the presence of a single spatial frequency channel for spatial frequencies below 0.6 cycle/degree and a separate channel for spatial frequencies between 0.6 and 1.0 cycle/degree. Waugh & Hess (1994) also examined the temporal sensitivity of the visual system using Gabor patches with a centre spatial frequency of 0.12, 1, and 3 cycles/degree. Using these stimuli, they assessed the contrast sensitivity function of the visual system at various temporal frequencies. Much like the results observed by Shoham et al., they observed that the peak of the temporal contrast sensitivity function shifted to lower frequency values as the spatial frequency of the Gabor patch was increased from 0.12 to 3 cycles/degree.

The presence of spatial frequency channels in stereopsis was recently demonstrated by Prince, Eagle, and Rogers (1998). Prince et al. examined the contrast masking function of test stimuli comprising a signal spatial frequency that varied between 0.4 to 3.2 cycles/degree in steps of 1 octave. These signal spatial frequencies were then combined with a binocularly uncorrelated noise mask with a centre spatial frequency varying between 0.4 to 9.5 cycles/degree in steps of one half octave. The idea behind such stimuli is that disparity acts as a signal which must be detected in the presence of the inherent noise of the visual system. If processing of disparity information is effected through separate spatial frequency channels, then one could potentially drown the disparity signal in noise by adding binocularly

uncorrelated noise with a spatial frequency at, or near, the spatial frequency of the stimulus whose disparity must be detected (i.e., the signal spatial frequency). The results observed by Prince et al. revealed masking functions which peaked at, or near, the signal spatial frequencies of the test stimulus (i.e., 0.4, 0.8, 1.6, and 3.2 cycles/degree). The interesting aspect of Prince et al.'s results is that they are in broad agreement with the results observed by Waugh and Hess (1994) for temporal frequencies.

Taken together, results from Prince et al. and Waugh and Hess indicate that processing of disparities is effected by channels which are tuned to spatial and temporal frequencies. Essentially, channels tuned to low spatial frequencies exhibit a high temporal resolution whereas channels tuned to higher spatial frequencies exhibit a progressively lower temporally-passed response as spatial frequency increases.

Although there is still some debate as to whether the spatiotemporal frequency preferences of the visual system at such low spatial frequencies represent a single channel clasping a continuous gradient of neurons tuned in space and time or separate narrowly tuned spatiotemporal channels (Issa, Trepel, & Striker, 2000), the fact that there are differences in the temporal frequency responses of the visual system across spatial frequencies puts into question whether the relatively high temporal resolution (30 to 40 Hz) observed by Morgan and Castet would hold beyond 1 cycle/degree. Forsooth, it is not unreasonable to suggest that the restricted low spatial frequency

range used in their experiment represents the upper limit of the temporal resolution of the visual system, and that we should observe poorer temporal responses for spatial frequencies above 0.32 cycle/degree. In summary, the marked differences in the spatiotemporal characteristics of the visual system for spatial frequencies between 0.12 and 3 cycles/degree (Shoham et al., 1997; Waugh & Hess, 1994; Prince et al., 1998) suggest that the consistent breakdown in performance between 30 and 40 Hz is likely to be the result of the relatively narrow range of low spatial frequencies used in the work of Morgan and Castet.

Another question is how the results obtained by Morgan and Castet can be generalised to more complex stimuli such as compound gratings comprising a high and a low spatial frequency. This is important as the natural environment is comprised of stimuli that are invariably more complex in their spatial frequency content. One aim of this study is to assess how the visual system handles the computation of disparity information for laterally drifting compound gratings. If indeed the upper temporal frequency limit of stereopsis decreases with increasing spatial frequency then it may be possible to create a compound grating with spatial frequency components which are processed with different temporal resolutions by the visual system. The question then arises as to which of the spatial frequencies would dominate the percept. One possibility is that the visual system's ability to compute disparities at high velocities is solely determined by the lowest spatial frequency

present in the stimulus (low spatial frequency constraint hypothesis). In which case we would expect the upper velocity limit of the compound grating to conform to that of the lower spatial frequency component when tested alone. Another possibility is that the presence of higher spatial frequencies will constrain the upper velocity limit as the visual system is more sluggish in processing higher spatial frequencies (high spatial frequency constraint hypothesis). As will be discussed in the following two sections, there is ample empirical evidence to expect either of these outcomes.

Low to High Spatial Frequency Constraints on Stereopsis

In this section, empirical evidence suggesting an important role of low spatial frequencies in resolving the correspondence problem in stereo vision is presented (Marr & Poggio, 1979; Legge & Gu, 1989; Wilson, Blake, & Halpern, 1991). As will be discussed, some authors (Marr & Poggio, 1979) suggested that the inherent ambiguity of stereo matches at higher spatial frequencies could be resolved if stereo matches at lower spatial frequencies were established prior to those at higher spatial frequencies. Empirical evidence for the importance of lower spatial frequencies in stereopsis was offered in a study conducted by Wilson, Blake, and Halpern (1991) who demonstrated that the presence of low spatial frequencies can actually constrain the fusion range at higher spatial frequencies. However, an influence of higher spatial frequencies on the fusion range of lower spatial frequencies was not observed. These results suggest that the visual system may have access to stereo matches at lower

spatial frequencies prior to, and independently from, stereo matches at higher spatial frequencies. If this is the case, the visual system may rely mainly on lower spatial frequencies when computing the depth of a drifting compound grating. As a result, we would expect little influence of the higher spatial frequency component of a compound grating on the upper velocity limit of stereopsis.

Initial demonstrations of stereopsis with random-dot stereograms (Julesz, 1960) have had a great influence on the formulation of models of stereopsis. Indeed, because of the fine spatial scale of the elements used in random-dot stereograms it was difficult to imagine how the visual system could achieve stereopsis with such stimuli without considerable computations, or trial and error. Yet stereopsis seemed to occur quasi-effortlessly for observers with normal binocular vision. Given the ambiguous nature of interocular correspondence for such stimuli and the ease with which they can be fused, most current models of stereopsis have focussed on the need for the visual system to resolve efficiently the massive amount of false matches that would result from attempting to fuse, on a point per point basis, two random dot stereograms.

There is empirical evidence of the inherent ambiguity that the visual system encounters when attempting to establish stereopsis at high spatial frequencies (e.g., Piggins, 1978; Odom & Chao, 1987; Legge & Gu, 1989). For instance, in an experiment conducted by Legge and Gu (1989), observers were asked to report

whether a test stimulus composed of two vertically oriented sine wave gratings (test stimulus) appeared in front of, or behind, a zero disparity grating which was presented just below the test stimulus. The threshold disparity for the detection of depth was then determined as a function of the spatial frequency of the grating. Results of the experiment showed that threshold disparity was inversely proportional to spatial frequency, reaching a minimum at about 2.5 cycles/degree, after which the threshold increased erratically. The irregular increase in disparity threshold above 2.5 cycles/degree can be explained by the fact that, when spatial frequency increases beyond a certain point, disparity becomes ambiguous since it is not clear which bar should be paired with which (classically referred to as the correspondence problem or more commonly the wall paper illusion).

With periodic stimuli (such as spatial frequency gratings), there is a geometrical reason why disparity detection should be linked to spatial frequency (Howard & Rogers, 1995). Indeed, the largest disparity over which two images of a grating can be matched is one half of the spatial period of the grating. If the disparity is any greater than this, the two gratings may be matched in more than one way. Because of this, only lower spatial frequencies allow the visual system to match two images at disparities exceeding the periodicity of higher spatial frequencies. In addition, because there are fewer contours at lower spatial frequencies, there are fewer chances of establishing erroneous matches. It is thus not surprising that many models

of stereopsis have utilized ideas and findings emanating from research on spatial frequency channels to explain how stereopsis is achieved. From as early as 1972, models such as that proposed by Felton, Richard, and Smith (1972) have suggested that disparities are processed by distinct size channels in the visual system, where small disparities are processed by high spatial frequency channels and large disparities by low spatial frequency channels.

Another model implementing spatial frequency channels in stereopsis was proposed by Marr and Poggio (1979). In an effort to offer a solution to the ambiguity of stereo matches with random-dot stereograms, Marr and Poggio (1979) suggested that the visual system could eliminate matching ambiguities by using disparity outputs at coarser scales to drive stereo matching at finer scales. They reasoned that it would be more efficient for the visual system to first find matches at low spatial frequencies and to use the disparity output of these channels to shift the range of stereo matching at higher spatial frequencies. There are a number of mechanisms which have been proposed to enable low spatial frequency channels to shift the matching process at higher spatial frequencies within the appropriate range of disparities. For example, in Marr and Poggio's (1979) model, this would be accomplished through initial vergence eye movements driven by initial disparity estimates from coarser scales. Regardless of the processes involved, these models all share the assumption that disparity estimates at larger scales are derived first and serve to guide stereo matches at finer scales.

Such an approach to the computation of stereoscopic depth suggests that the visual system makes use of the spatial frequency content of stimuli to compute disparities. In fact, there is much support in the psychophysical literature for a disparity encoding scheme based on spatial frequencies. As stated earlier, from a theoretical perspective, with periodic stimuli such as spatial frequency gratings, there are geometrical reasons why disparity detection must be linked to spatial frequency (Howard & Rogers, 1995). Indeed, the largest disparity over which two images of a grating can be matched in a particular way is one half of the spatial period of the grating (i.e. $\pm 180^\circ$). Based on a disparity encoding scheme derived from spatial scales, we would expect that both the smallest and the largest disparity over which two gratings of identical spatial frequencies can be matched would be related in some way to a constant angular phase disparity.

The threshold data obtained by Legge and Gu for spatial frequencies up to 2.5 cycles/degree are consistent with the idea that disparity thresholds are best described as a constant angular phase offset. Indeed, their results demonstrated that the log/log plot of the disparity threshold in arc sec was inversely proportional to spatial frequency, reaching a minimum at about 2.5 cycles/degree. The proportional relationship between the decrease in the threshold, as expressed in arc sec, with increasing spatial frequency indicates that the visual system uses spatial frequency phase offsets as a means to compute binocular disparities. In fact, converting arc sec

threshold as an angular phase disparity results in a relatively constant threshold angular phase disparity for all spatial frequencies. Similar results were also observed in an experiment conducted by Smallman and Macleod (1994) where they examined the upper disparity limit of stereopsis. Results of their experiment indicated that the fusion limit of human observers decreases with increasing spatial frequency in a manner that is, again, consistent with the prediction of a phase encoding scheme up to a spatial frequency of 2.5 cycles/degree. The proportional reduction in both the disparity threshold and the upper disparity limit with increasing spatial frequency is referred to as the size-disparity correlation.

The idea that low spatial frequencies play a pivotal role in stereopsis has also received empirical support in a series of experiments conducted by Wilson, Blake, and Halpern (1991). In their experiment, Wilson and his colleagues measured the effects of a background grating of one spatial frequency on the fusion limit of another grating. The question addressed was whether the depth plane defined by a low spatial frequency grating had any effect on the fusion range of a high spatial frequency patch (8 or 12 cycles/degree) (as implied by Marr & Poggio, 1979). To answer this question, the fusion range of the high spatial frequency patch was measured in the presence and in the absence of the low spatial frequency grating.

Results of their experiments showed that low spatial frequencies restrict the fusion range of higher spatial frequencies. However, the influence of lower spatial

frequencies was only observed when the low and high spatial frequencies were separated by no more than 2 octaves. When they were separated by more than 2 octaves, no interactions across spatial frequencies were observed and the high frequency patch was seen transparently through the low spatial frequency grating. More importantly, no evidence was found for the case where higher spatial frequencies constrained the fusion range of lower spatial frequencies. Based on these results, Wilson et al. concluded that low spatial frequencies constrained the processing of disparities at higher spatial frequencies but only within a range of two octaves.

SUMMARY

In this section, ecological reasons as well as empirical evidence for a low spatial frequency constraint on stereopsis were discussed. From an ecological perspective, one can readily see why the visual system would employ a low to high spatial frequency coding scheme for disparity. By first matching stimuli on low spatial frequencies where ambiguity is less likely to be a problem and then using these coarse disparity estimates to drive stereo matches at higher spatial frequencies, the visual system could avoid many of the mismatches that arise at higher spatial frequencies where interocular correspondence is more ambiguous.

In addition to this, empirical evidence was reviewed which suggests that a low to high spatial frequency constraint is at work in the processing of disparity. These studies have shown, in turn, that there is an inherent ambiguity in stereopsis associated

with higher spatial frequencies (Legge & Gu, 1989) and that, consistent, with these data, the visual system seems to prioritize binocular matches at lower spatial frequencies when processing disparities (Wilson et al., 1991). This, presumably, allows the visual system to then use coarse disparity estimates at lower spatial frequencies to guide stereo matches at higher spatial frequencies. The idea that the visual system can only constrain stereo matches at higher spatial frequencies after having derived the appropriate disparity signal at lower spatial frequencies, strongly suggests that disparity signals at lower spatial frequencies are derived independently from, and prior to, disparity signals at higher spatial frequencies.

In considering the question that lead to this review of the literature (i.e., upper velocity limit of stereopsis for compound gratings), there are good reasons to believe that the upper velocity limit of stereopsis for a compound grating would be determined by the lower spatial frequency component. This is despite the lower temporal resolution of the visual system at higher spatial frequencies. Indeed, intuitively, one would expect that if the visual system gains independent access to disparity information at low spatial frequencies where temporal resolution is optimal, then such a signal would be available when processing the disparity of a laterally moving object.

High to Low Spatial Frequency Constraints on Stereopsis

In this section, we present a number of studies which demonstrated that high spatial frequencies may also influence stereo matches at lower spatial frequencies. Overall, three sources of evidence are presented which suggest that, in stereopsis, high spatial frequencies interact strongly with lower spatial frequencies. Indeed, the resistance of higher spatial frequencies to factors which are known to be detrimental to stereopsis at lower spatial frequencies (Cormack, Stevenson, & Landers, 1991), the domination of higher spatial frequencies at small disparities (Rohaly & Wilson, 1993), and the presence of high spatial frequency constraints in motion (Clear & Braddick, 1990) argue strongly for the presence of high spatial frequency constraints on the upper velocity limit of stereopsis. That is to say, the visual system may not have independent access to the lower spatial frequency component in drifting compound gratings.

Stereopsis for static Images

Although there is much evidence to suggest that processing of disparities proceeds from coarse to fine spatial scales, there is also evidence that disparities at high spatial frequencies can serve to disambiguate disparities at lower spatial frequencies. For instance, in an experiment conducted by Smallman (1995), observers were asked to view simultaneously an ambiguous and an unambiguous stimulus specified at different spatial frequencies. The ambiguous stimulus could be matched

in the crossed or uncrossed directions equally well. It was created by presenting a low spatial frequency grating vertically. To achieve ambiguity, the grating was presented 180 degrees out of phase across the two eyes with an odd number of half cycles presented to the left eye and two less to the right eye. This produced a fused stimulus that was completely ambiguous and which could be matched in the crossed or uncrossed direction equally well. The unambiguous stimulus consisted of rectangularly filtered random-dot stereograms, shifted inside panels with zero disparity. As noted by Smallman, rectangularly filtered stereograms specify disparities that are unambiguous because the effective horizontal frequency varies vertically as well as horizontally, thus avoiding the wallpaper illusion inherent with vertically modulated sine waves.

On each trial, observers were asked to view two panels (i.e., test and standard panels). The test panel was presented 15 arc min above the fixation point and contained filtered stereograms at a given disparity and with a centre spatial frequency of either 2 or 8 cycles/degree. To these stereograms several cycles of the ambiguous grating with a spatial frequency of 2 or 8 cycles/degree were added. Thus, the display could comprise a low spatial frequency ambiguous stimulus with a high spatial frequency unambiguous stimulus or vice-versa. The depth of this panel was judged against the standard panel situated 15 arc min below the fixation point and at zero disparity. The standard was made of 2 cycles/degree filtered stereograms of the same

type as that used for the unambiguous test stimulus. The question asked by Smallman was whether the ambiguous low frequency grating would be seen in front of the fixation plan because the random-dot stereogram also lay in front, or whether it would be seen randomly in front or behind. If it were reliably seen in front, it could be concluded that the high frequency pattern had influenced the stereo matching of the low spatial frequency grating.

Contrary to the strict low spatial frequency constraint hypothesis, results from the fine to coarse condition (i.e., ambiguous low spatial frequency grating coupled with an unambiguous high spatial frequency stereograms) demonstrated that disambiguation was significantly greater than what would be expected by chance alone (i.e., 50%). In fact, the disambiguation of the low spatial frequency grating by the high spatial frequency pattern was observed for a variety of disparities peaking at disambiguation rates of 87% and 98%. Although the disambiguation was not as great as that observed in the coarse to fine direction, the fact that there was disambiguation demonstrates that higher spatial frequencies can influence stereo processing at lower spatial frequencies. In other words, the processing of disparities does not always have to proceed from coarse to fine.

There is a stronger line of evidence in support of a high to low spatial frequency constraint hypothesis. A number of studies (Kaufman, Bacon, & Barosso, 1973; Foley, 1976; Parker & Yang, 1989; Rohaly & Wilson, 1993) have reported a

phenomenon commonly referred to as disparity averaging or depth pooling. Disparity averaging is the phenomenon by which two stimuli with different but similar disparities are perceived as a single stimulus appearing at an intermediate depth. Experiments in depth averaging have been widely used to estimate the tuning of disparity channels as well as assessing the extent to which different spatial scales interact in the range of disparity differences over which disparity averaging occurs. In a study conducted by Rohaly & Wilson (1993), observers were asked to view a dichoptic pair of spatial frequency gratings that were superimposed on a second pair with a different spatial frequency. One of the two pairs (either the low or high spatial frequency grating) was presented at zero disparity while the disparity of the other pair was varied. Observers were then asked to compare the perceived depth of this grating with that of an identical reference display in which the disparity of both the low and high spatial frequency component was equal.

Their results revealed that, for as long as the spatial frequencies used in the test grating did not differ by more than 3 octaves, the test grating was perceived at intermediate depths over disparity differences of up to 112.3 arc sec. More interesting, however, is the observation that the perceived intermediate depth plane was systematically biased towards the disparity carried by the higher spatial frequency. That is, when the higher spatial frequency carried the zero disparity signal, intermediate depth was perceived as lying closer to the zero-disparity plane. This

contrasts with the case where the higher spatial frequency carried the disparity signal, as the depth plane was perceived as lying closer to the actual disparity of the higher spatial frequency.

As noted by Rohaly and Wilson, the fact that higher spatial frequencies dominate the determination of an averaged depth plane seem to be in strict contradiction to the results observed by Wilson et al. who demonstrated that the presence of a disparity signal at lower spatial frequencies dictated the spatial extent over which a higher spatial frequency grating could be fused. In order to explain this discrepancy, Rohaly and Wilson noted that the range of disparities used in both experiments were quite different. Whereas Wilson et al. used mainly large disparities, Rohaly and Wilson restricted their observations to disparities at or below 112.3 arc sec. Based on this, Rohaly and Wilson argued that large disparities may favour disparity signals at low spatial frequencies where they constrain the fusion limit of high spatial frequencies, and that small disparities favour high spatial frequencies in determining the appropriate depth plane. This interpretation is quite consistent with the observation of a size-disparity correlation in disparity threshold and fusion-limit data (Legge & Gu, 1989; Schor, Wood, & Ogawa, 1984).

Yet another line of evidence supporting an important role for higher spatial frequencies in stereopsis was provided by Cormack, Stevenson, & Landers (1997). It is well known that mismatches in contrast across the two eyes are detrimental to

stereo matching. So detrimental that if the contrast in only one eye is reduced, disparity thresholds increase much more rapidly than if the contrast is reduced by the same amount in both eyes (Halpern & Blake, 1988; Legge & Gu, 1989; Schor & Heckman, 1989). This may not seem particularly strange until one realises that it implies that if the contrast in one eye is increased, stereo acuity is degraded. This is a phenomenon which Cormack, Stevenson, and Landers (1997) have termed the “stereo-contrast paradox”.

In a previous experiment, Cormack et al. (1996) reported that the stereo-contrast paradox did not occur at higher spatial frequencies. That is, stereo matches made at higher spatial frequencies were relatively immune to unequal variations in contrast across the two eyes. In their experiment, Cormack et al. (1996) investigated the effects of interocular differences in contrast with the use of dynamic random element stereograms (DRESs) instead of spatial frequency gratings, such as those used in previous experiments (Schor & Heckman, 1989; Legge & Gu, 1989). In their first experiment using DRESs, they failed altogether to show a stereo-contrast paradox effect. Indeed, their results demonstrated that disparity thresholds in DRESs were limited to the lower of the two image contrasts instead of the accelerated loss observed by Legge and Gu. In other words, reducing the contrast in one eye only had the same effect as reducing the contrast in both eyes by the same amount.

To explain the discrepancy between their results and those observed by Legge and Gu, Cormack and his colleagues (1997) later suggested that the reason they failed to observe a stereo-contrast paradox was because of the high spatial frequency content of their stimuli and that, somehow, this allowed the visual system to overcome the difference in inter-ocular contrast which had proved so detrimental to stereopsis at lower spatial frequencies. To support this idea, Cormack et al. (1997) remarked that most studies that reported a stereo-contrast paradox used low spatial frequency stimuli (Simmons, 1984; Legge & Gu, 1989; Schor & Heckman, 1989), whereas studies which used a wider range of spatial frequencies generally reported a reduction or an absence of the effect (e.g., Halpern & Blake, 1988).

To test this hypothesis, Cormack et al. (1997) measured the disparity threshold of observers to spatial frequency gratings of 1, 3, and 5 cycles/degree, with contrast ratios of 8:1, 4:1, 2:1, and 1:1. Their results demonstrated an accelerated loss in stereo acuity at low spatial frequencies with every increase in the inter-ocular contrast ratio, thus replicating the results of Legge and Gu and Schor and Heckman (1989). As the spatial frequency of the grating was increased, however, the magnitude of the stereo-contrast paradox decreased proportionately up to a point where disparity thresholds were solely determined by the lower of the two contrasts. That is, the disparity threshold function for unequal contrast was identical to that observed when contrast was reduced equally in both eyes. The fact that the visual system can tolerate

mismatches in contrast at high spatial frequencies but is susceptible to mismatches at lower spatial frequencies suggests that processing of disparities at high spatial frequencies could still be possible in the absence of a reliable disparity signal at lower spatial frequencies. Hence, although low spatial frequencies seem to play an important role in stereo matches, it is unlikely that the visual system relies solely on the disparity output at lower spatial frequencies to drive stereo matches at higher spatial frequencies.

Disparity and Motion

Another interesting demonstration of how high spatial frequencies can influence the processing of lower spatial frequencies comes from research in motion perception using random-dot kinematograms (RDKs). A RDK is composed of two identical but spatially displaced versions of a random-dot pattern composed of black and white dots. Generally, both frames are viewed through a stationary window so that the boundaries of the pattern are not visible. The results generally observed with these types of stimuli is that while for small jumps (i.e., displacements), observers report smooth and coherent motion (i.e., global motion), as the displacement between frames is increased the percept quickly becomes one of incoherent local motions. The displacement value at which the direction of the global motion can be reliably reported is termed d_{max} .

One of the most intriguing results observed in studies that have used RDKs suggests that the perception of motion is restricted by the high spatial frequency content of visual stimuli (Cleary & Braddick, 1990; Eagle, 1998). In an experiment conducted by Cleary and Braddick (1990), d_{max} values of observers were measured using low-pass filtered RDKs. The results obtained by Cleary and Braddick showed that spatial filtering of higher spatial frequencies caused d_{max} values to increase in proportion to the extent of the filtering. These results are indicative of the visual system's inability to disregard higher spatial frequencies in broadband stimuli. To explain these results, Cleary and Braddick suggested that the presence of high spatial frequencies served to constrain d_{max} and that filtering the high spatial frequencies enabled the visual system to rely on lower spatial frequencies which support larger d_{max} values.

High spatial frequency constraints on d_{max} were also observed by Eagle (1998) who demonstrated that constraints imposed by the presence of higher spatial frequency noise on d_{max} values were greater if they were within one octave of the signal spatial frequency. In his experiment, Eagle (1998) asked observers to view stimuli with a low signal spatial frequency spanning one octave (0.33 - 0.67 cycle/degree). High spatial frequency noise could then be added to the adjacent octave (0.67 - 1.33 cycles/degree) or to the next three non-adjacent octaves (1.33- 10.67 cycles/degree) which Eagle (1998) termed the "notch" stimulus because of the

one octave separation between the signal and the noise. Interestingly, despite the presence of greater noise in the notch stimulus, reductions in d_{max} values were greater in the adjacent noise condition. These results are similar to the findings of Chang and Julesz (1983) who found that removing the mid-range of spatial frequencies (i.e., between 3 and 6 cycles/degree) in RDKs had the same effect on d_{max} as when they are low-passed only.

Taken together, results from Eagle (1998) and Chang and Julesz suggest that introducing a separation between the signal and noise spatial frequencies allows the visual system to access independently the motion signal from either the low or the high spatial frequencies. Hence, it would seem that, much like interactions across spatial frequencies observed in stereopsis (Wilson et al, 1991; Julesz & Miller, 1975), interactions between spatial frequencies in motion occur mainly when they are within a given range of one another.

Cleary and Braddick's and Eagle's (1998) results gain further pertinence due to the increasing number of studies in psychophysics (Regan & Beverly, 1973; Anstis & Harris, 1974) and in neurophysiology (Maunsell & Van Essen, 1983; Bradley, Qian, & Anderson, 1995) that demonstrate that motion and disparity processing may share a common neural substrate. For example, Bradley et al. (1995) demonstrated that approximately 66% of sampled individual neurons in area MT exhibited selectivity to both direction of motion and disparity. Psychophysical studies have also shown that

the physiological properties of neurons that are selective for both disparity and motion are reflected at the behavioural level as well. Indeed, motion transparency, motion aftereffects, as well as motion adaptation have all been shown to be disparity contingent (Regan & Beverly, 1973; Qian, Anderson, & Adelson, 1994; Kwas, Von Grunau, & Dubé, 1995).

In their experiment, Bradley et al. used stimuli consisting of preferred and anti-preferred random-dot patterns. Preferred random-dot patterns comprised a motion and a disparity signal in the preferred direction and depth of the tested neuron, whereas the anti-preferred pattern comprised a motion signal in the opposite direction of that preferred by the neuron and a disparity signal which could assume a variety of depths. The preferred and anti-preferred patterns were superimposed to form a transparent motion stimulus with opposite motion across patterns. These stimuli were created in order to assess a well known phenomenon called motion-opponent inhibition where motion signals in opposite direction tend to inhibit each other within a given MT receptive field (Snowdon, Treue, Erickson, & Andersen, 1991). By varying the disparity signal of one of the patterns, Bradley et al. wanted to assess whether this inhibition would be present when patterns with opposite direction of motion are presented at different depth planes. Results from their experiment showed that in 66% of MT neurons tested, inhibition due to the presence of opposite direction of motion was only present when the two patterns were presented at or near the same

depth plane. These results suggest that disparity information plays a critical role in determining which motion signals should be processed jointly and which should be processed independently.

That disparity plays an important role in binding motion signals has also been demonstrated psychophysically by Kwas, Von Grunau, and Dubé (1995) in a motion adaptation study. It is now well known that prolonged viewing of moving plaid stimuli yields bistable percepts, where perception switches between separate gratings appearing to slide over each other (component motion) and a percept of coherent motion moving in a third direction (pattern motion). Von Grunau and Dubé (1993) had previously shown that unambiguous adaptation to either component motion or pattern motion can reduce the proportion of time in which component or pattern motion is perceived in a subsequently presented ambiguous pattern. Accordingly, adaptation to unambiguous pattern motion decreased the time during which pattern motion was perceived and adaptation to unambiguous component motion reduced the time during which component motion was perceived.

One interesting aspect of plaid stimuli is that binocular disparities can be introduced in either grating comprising the plaid independently. Thus, a plaid could be generated where one of its components is presented at zero disparity and the other at a different depth plane. Kwas et al. asked observers to adapt to plaids containing either an unambiguous pattern motion or an unambiguous component motion signal.

The components of the adaptation plaid were always presented at zero disparity. The test stimulus comprised ambiguous motion with components that could be presented on different depth planes. Their results demonstrated that presenting the gratings comprising the test plaid on different depth planes reduced the time during which pattern motion was perceived regardless of the adapting condition (i.e., unambiguous pattern or component motion).

Overall, the results of studies that have examined spatial frequency interactions in motion (Cleary & Braddick, 1990; Eagle, 1998; Chang & Julesz, 1983; Bradley et al., 1995; Kwas et al., 1995) suggest that both distance in spatial frequency and distance between depth planes can induce independent processing of motion signals. This is sensible as motion signals arising from a pair of stimuli with identical or adjacent spatial scales presented on a common depth plane may indicate to the visual system that these motion signals are arising from a common object. An efficient response would suggest that the visual system bind the motion signals arising from these stimuli as a single motion signal. In view of these studies, insofar as the upper velocity limit of stereopsis is determined by the temporal resolution of neurons which are selective for both disparity and direction of motion, one would expect that the high spatial frequency constraints observed in motion (Cleary & Braddick, 1990; Eagle, 1998) would also be observed in tasks requiring observers to detect disparity signals in laterally moving stimuli.

Indirect evidence of common processing boundaries (spatial and temporal) for motion and stereo vision was provided by Glennerster (1998). Using a d_{max} paradigm, Glennerster recently demonstrated that, in the binocular domain, stereopsis and motion share similar limitations in solving the correspondence problem. As argued by Glennerster stereopsis and motion both require the detection of correlation between images. For motion, the correspondence problem can be described as correlating images in time, whereas for stereopsis, the correspondence problem can be described as correlating images in space. In his experiment, Glennerster presented a pair of random-dot patterns. Across patterns, dots in one of the frames were displaced to the right or to the left. To generate motion, the images were displayed as a binocular two-frame apparent motion sequence with each frame being presented for 150 msec. Thus, for these displays, observers were correlating the displacement in time as they would for motion. For stereo, the two images were presented as a binocular pair for 150 msec. Thus, for these displays, observers were correlating displacement in space as they would for stereopsis.

Glennerster's results indicated that for a wide range of conditions, stereo and motion yielded nearly identical d_{max} values. Although these results do not provide direct evidence of a common neural substrate for the processing of motion and stereopsis, in the context of recent physiological studies that point to a common neural

substrate, it is interesting that stereo and motion would share common limitations in solving the correspondence problem.

SUMMARY

Stereo acuity for laterally moving objects is an unexplored area of stereopsis. In fact, an extensive review of the literature in both the fields of psychophysics and neurophysiology has uncovered only one study that has specifically examined this issue (Morgan & Castet, 1995). The results obtained by Morgan and Castet suggest that the temporal resolution of the visual system for processing disparity in lateral motion is approximately 30 to 40 Hz regardless of spatial frequency. However, based on a review of the literature, how the visual system will establish stereopsis for laterally moving gratings remains largely an empirical question. Indeed, although the study conducted by Morgan and Castet provides insights into the temporal resolution of the visual system for disparity processing in dynamic displays, it is clear that the range of spatial frequencies used in their study affords only a partial perspective with respect to the upper temporal frequency limit of stereopsis. Results obtained by Shoham et al., Waugh and Hess as well as Prince et al., suggest that, for laterally moving stimuli, disparity signals at higher spatial frequencies could not be maintained at such high temporal frequencies.

The issue is further complicated with the introduction of compound gratings. This is because a compound grating could potentially be constructed with spatial

frequencies that are processed with differing temporal resolution by the visual system. Recognizing this, it is not obvious how the visual system would respond to such stimuli. The inherent ambiguity of stereo matches at higher spatial frequencies suggests that stereopsis should proceed by first establishing matches at lower spatial frequencies to then use these initial disparity estimates to drive stereo matches at higher spatial frequencies (Marr & Poggio, 1979; Legge & Gu, 1989; Wilson et al, 1991). This suggests that disparity signals from lower spatial frequencies are available prior to and independently from disparity signals at higher spatial frequencies. An efficient visual response to a drifting compound grating would suggest that the visual system would rely on the disparity output of lower spatial frequencies where the visual system is optimally tuned to high temporal frequencies. However, the resistance of higher spatial frequencies to factors which are known to obliterate stereopsis at lower spatial frequencies (Cormack et al., 1997), the ability of higher spatial frequencies to disambiguate stimuli at lower spatial frequencies (Smallman, 1995), as well as the domination of higher spatial frequencies in depth averaging (Rohaly & Wilson, 1993), argue for a less passive role of higher spatial frequencies in stereopsis.

Finally, the absence of interactions across spatial frequencies in motion (Chang & Julesz, 1983; Eagle, 1998) and stereopsis (Wilson et al., 1991; Bradley et al., 1995) with a 2 octave separation, suggest that high spatial frequency constraints on the

processing of disparity in lateral motion should be limited to compound gratings comprising adjacent spatial frequencies (within a one octave separation). Thus, for compound gratings with spatial frequencies separated by 2 octaves or more, we would expect the upper velocity limit to conform to that of the lower spatial frequency.

In this thesis, stereo performance at high velocities was examined over a wider range of spatial frequencies (0.08 to 2.56 cycles/degree). As suggested by previous studies, it was expected that the temporal frequency limit for stereopsis could not be maintained at 30 to 40 Hz beyond 0.32 cycle/degree. The upper velocity limit of stereopsis performance for compound gratings with a 1 or 2 octave separation was also assessed. Finally, we examined the effect of disparity magnitude on the upper velocity limit of stereopsis for compound gratings with a 1 octave separation².

Overall it was found that : (1) the upper temporal frequency limit of stereopsis is dependent on spatial frequency in such a way that as spatial frequency is increased beyond 0.32 cycle/degree, the temporal frequency at which depth can no longer be reliably reported decreases. (2) Stereo performance for compound gratings shows a

2

It should be noted that velocity (deg sec⁻¹) and temporal frequency (Hz) represent different ways of expressing motion. Velocity, however, is more convenient to use for compound gratings as spatial frequencies within a compound grating share a common velocity, but present differing temporal frequencies. For simple gratings, however, expressing performance as a function of temporal frequency is more convenient. This is because, if channels tuned to different spatial frequencies possess equivalent temporal resolutions, this would be readily observed on the temporal frequency dimension as equal upper temporal frequency limits. When expressed as a velocity, equivalent temporal resolutions with increasing spatial frequency would result in a proportional reduction in the upper velocity limit. Thus, when dealing with simple gratings, expressing stereo performance as a function of temporal frequency is preferable. In this thesis, results are expressed as a velocity when dealing with compound gratings. For simple gratings, however, results are expressed both as a temporal frequency and as a velocity.

high spatial frequency constraint at small disparities which (3) are largely eliminated at larger disparities. These results are consistent with the modern emerging view of stereo vision which is characterised by the multiplexing of motion and disparity.

EXPERIMENT I

Morgan and Castet examined the upper temporal frequency limit of stereopsis for simple gratings with spatial frequencies ranging between 0.04 and 0.32 cycle/degree. They observed a constant upper temporal frequency of 30 to 40 Hz. An important premise of this study is that the upper temporal frequency limit will fall below 30 to 40 Hz for spatial frequencies above 0.32 cycle/degree. In Experiment I, the upper temporal frequency limit of stereopsis was determined over a wider range of spatial frequencies (0.08, 0.16, 0.32, 0.64, 1.28, and 2.56 cycles/degree). Four observers were tested at all spatial frequencies except at 0.08 cycle/degree where only two of the four observers were tested. The upper range of spatial frequencies was limited to 2.56 cycles/degree in order to avoid the inherent spatial ambiguity (i.e., wall paper illusion) associated with higher spatial frequencies (Legge & Gu, 1989). However, as reported by Shoham et al.. and Waugh and Hess, the temporal resolution of the visual system for spatial frequencies differs even within this range. Thus, differences in the upper temporal frequency of stereopsis should be observed within the 0.08 to 2.56 cycles/degree range. If, as suspected, the disparity signal at higher spatial frequencies can not be maintained at the same temporal frequencies as those

observed by the Morgan and Castet, a gradual reduction in the temporal frequency at which depth can no longer be reliably reported should be observed as spatial frequency is increased beyond 0.32 cycle/deg.

METHOD

Observers: Four observers served in this experiment. All observers had 20/20 or corrected to 20/20 vision. Of the four observers, three were naive as to the purpose of the experiment.

Apparatus: Stimulus displays were presented on a Hitachi monitor driven by a Cambridge Research Systems VSG 2/3 video board mounted in an IBM compatible computer. The vertical scan rate of the monitor was set at 150Hz or one frame every 6.7 msec. The horizontal and vertical extents of the display comprised 928 pixels and 401 lines respectively. The computer monitor was divided horizontally into two equal extents of 16 cm. From a distance of 75 cm, the left and right half of the monitor subtended approximately 12.0 degrees of visual angle with each pixel subtending 0.025 degree of visual angle or approximately 1.5 arc min.

The monitor was calibrated using the CRS OptiCal system to achieve a linear gray scale centred on a mean luminance of 49.7 cd/m² (peak and trough luminance were 88.9 and 3.2 cd/m² respectively). Intensity values generated by the look up table are presented in figure 1. As can be observed, aside from a slight non-linearity in the first 5 intensity values, the monitor yielded a fairly linear gray scale. Stereoscopic

Figure 1

Figure 1 : Display intensity values (cd/m²) as function of look up table index

separation to the left and right eye was achieved with the use of polarized sheets attached to the monitor and two smaller sheets attached to the head rest. The head rest was equipped with prisms which aided convergence of the eyes to afford comfortable viewing. The total light transmittance of the system including the two sets of polarized sheets and the prisms was 10.5%. This resulted in a peak and trough luminance of 9.25 and 0.34 cd/m^2 respectively. This allowed the generation of spatial frequency gratings with a Michelson contrast of approximately 92.9%. Crosstalk across dichoptic separation of the monitor was assessed by measuring the luminance transmittance of the polarised sheets with their polarised orientation in quadrature. The total transmittance due to crosstalk was approximately 0.004% with a peak luminance measured at 0.014 cd/m^2 .

Stimuli and Procedure: The stimuli used in the experiment consisted of spatial frequency gratings of 0.08, 0.16, 0.32, 0.64, 1.28, and 2.56 cycles/degree. The presentation display consisted in a pair of identical gratings which were presented on the left and right half of the monitor. Within each half of the monitor, the vertical extent was divided into three equal vertical portions with the top and bottom thirds containing stationary static random-dot patterns presented at zero disparity. Although the top and bottom random-dot patterns within each horizontal half differed from each other, the top left and top right as well as the bottom left and bottom right patterns were perfectly correlated to ensure proper binocular fusion of each stereo half.

The vertical middle portions of the display contained the spatial frequency gratings to be fused. Figure 2 illustrates an example of the experimental displays used in this thesis. Disparities in the gratings were introduced by advancing and retarding symmetrically the right and left grating for crossed disparities and vice versa for uncrossed disparities (for a 1 pixel shift, the phase of the right or left grating was randomly shifted to the left or the right as needed). Care was taken in the displays so as to minimise monocular depth cues during presentation of dynamic displays. This was achieved by ensuring that as gratings are displaced on the monitor, pixels displaced off the monitor were replotted on the opposite side of the displaced image (i.e., images were wrapped around). We assessed the ability of observers to identify the sign of the disparity using static versions of the displays when viewed monocularly. Overall, results of this control condition indicate that monocular cues could only be successfully used at disparities far above (approximately 2.4 times greater) those observed during binocular viewing.

Before each trial began, observers were asked to fuse two crosses appearing in the centre of the right and left half of the monitor. This was done in order to ensure that observers converged their eyes appropriately before initiating a trial. When observers were certain that they perceived a single cross, they initiated a trial by pressing the bottom button of a three button response box. On each trial, the grating

Figure 2

Figure 2 : Sample experimental display depicting simple grating used in Experiment I.

was presented for 200 msec, preceded and followed by a mask consisting of a checker board pattern which filled the display area for 75 msec.

The task of the observer was to determine whether the fused grating appeared in front of, or behind, the zero disparity random-dot pattern (i.e., a 2AFC procedure). Observers responded to each trial by pressing the upmost button of the response box to indicate that the grating had appeared behind, or the middle button to indicate that the grating had appeared in front of the random-dot pattern. Observers were given no time limits between trials to respond, and correct responses were signalled with an auditory signal.

The threshold disparity for each observer was determined by the amount of phase shift (in pixels) needed to achieve a 75% correct performance using a static version of the display. Static disparity thresholds yielding 75% correct performance were estimated for each observer at each spatial frequency using a Parameter Estimation by Sequential Testing (PEST) procedure (Taylor & Creelman, 1967). After the 75% correct static disparity thresholds were estimated, observers were presented with dynamic displays in which the grating was displaced horizontally by the required number of pixels at every frame or multiples of frames to generate the intended velocity. The range of temporal frequencies tested was 5 to 40 Hz in steps of 5 Hz. Due to the proportional relationship between velocity in deg sec^{-1} and spatial frequency, the range of temporal frequencies to be tested generated velocities which

could range between 62.5 and 500 deg sec⁻¹ for a spatial frequency of 0.08 cycle/degree, and between 2 to 15.6 deg sec⁻¹ for a spatial frequency of 2.56 cycles/degree. On each trial, the direction of motion (left or right) and disparity (crossed or uncrossed) was randomly determined. Spatial frequency gratings were tested in separate sessions. Two observers were tested with increasing spatial frequencies beginning with the lowest whereas the other two observers were tested with decreasing spatial frequencies beginning with the highest. Within each session, stereo performance for each velocity was tested in separate blocks of trials of increasing velocity until the estimated disparity threshold was doubled that of the static condition. Observers were then tested at a velocity intermediate to the last two highest tested velocities. Observers were allowed to take breaks as they saw fit. However, a mandatory 10 minute break was imposed after each session.

Observers were trained in a number of sessions prior to commencing experimental sessions. This was done to ensure that performance during experimentation was devoid of practice effects. Practice sessions were generally carried out over a period of 3 to 4 weeks and typically included two- to three-hour sessions twice a week. Experimental sessions began when observers demonstrated no more improvement in performance over three consecutive sessions. In addition to this, each experimental session began with a brief practice session using static versions of the displays at the spatial frequency tested during the session. A total of 10

sessions per spatial frequency were conducted, yielding 10 disparity threshold estimates per point.

RESULTS

Stereo Thresholds for Static Displays

Because relative disparity thresholds for drifting gratings are established with respect to the disparity threshold of static displays, it is important to ensure that the results observed in the static condition are consistent with previous studies (Legge & Gu, 1989; Smallman & Macleod, 1994). This is to ensure that comparisons made between disparity thresholds for drifting and static gratings are based on expected results in the static condition.

Figure 3 depicts static disparity thresholds (arc min) for each observer as a function of spatial frequency. As can be observed, results across the four observers describe a gradual decrease in the disparity threshold as spatial frequency increases. These results are consistent with those observed by Legge and Gu as well as Smallman and Macleod (1994) who demonstrated that, within the range of spatial frequencies used in this experiment (i.e., 2.5 cycles/degree or less), threshold disparity in arc min, is inversely proportional to spatial frequency. When the data are converted as phase angle disparities, thresholds remain relatively constant across spatial frequency with results varying between 20 and 25° degrees of phase angle offset owing to slight individual differences. Again, these results suggest that within this range of spatial

Figure 3

Figure 3 : Static disparity thresholds (arc min) as a function of spatial frequency. Error bars represent two standard errors

frequencies, disparity encoding is best described as a phase angle offset (Legge & Gu, 1989; Smallman & MacLeod, 1994; Freeman & Ohzawa, 1990; DeAngelis, Ohzawa, & Freeman, 1991; 1995).

Stereo Performance for Drifting Gratings

Panels a, b, c, and d in figure 4, depict stereo performances for observers SL, NC, LR, and JG across spatial frequencies as a function of velocity. The Y-axis on the figure depicts stereo performance for drifting gratings as a relative phase angle disparity threshold. The relative disparity threshold was derived by comparing the phase angle disparity needed to achieve a 75% correct performance at each velocity with that of the appropriate static display. Hence, the Y-axis is interpreted as the difference between threshold disparities using dynamic and static displays (i.e., dynamic disparity threshold minus the static disparity threshold).

As can be observed, for all observers, there is a gradual decrease in the velocity at which depth can no longer be reliably reported as spatial frequency is increased. These data are consistent with the results obtained by Morgan and Castet and depict a gradual decrease in the upper velocity limit of stereopsis as a function of spatial frequency. A two parameter function describing an exponential growth in the relative disparity threshold as a function of velocity was fitted to the data (see equation 1) and is shown to provide a good fit to the results with an average adjusted R squared value of .967.

Figures 4 a and b

Figures 4 a and b : Relative disparity thresholds (simple gratings) as a function of velocity

Figures 4 c and d

Figures 4 c and d : Relative disparity thresholds (simple gratings) as a function of velocity

$$\text{rdt} = \exp^{\alpha(v-\chi)} \quad \text{Equation 1.}$$

Where rdt = relative disparity threshold, v is the velocity (in deg sec^{-1}), and α and χ are constants, representing the rate of growth and the point of acceleration of the function respectively.

One feature of this function is that the slope of the function is at a 45 degree angle when velocity equals χ on a linear/linear plot. The χ parameter thus provides us with an estimate of the “knee” of the function which can be interpreted as the point at which the disparity needed to yield a 75% correct response rate departs from the static disparity threshold (i.e., upper velocity limit).

Using the fitted χ parameter, estimated upper velocity limits in stereo performance were obtained for each observer at each spatial frequency. These estimated upper velocity limits were then converted as temporal frequencies by multiplying its value with the appropriate spatial frequency (see appendix I for individual data). As expected, the range of estimated upper velocity limits in stereo performance differs markedly depending on the spatial frequency of the grating. Across observers, stereo performance can be maintained at very high velocities with values ranging between 290 and 412 deg sec^{-1} for the 0.08 cycle/degree grating.

whereas the range of performance for the highest spatial frequency (2.56 cycles/degree) was limited to values ranging between 0.2 and 6.5 deg sec⁻¹ (see appendix 1 for individual data).

Figure 5 depicts the obtained estimated upper velocity limits (deg sec⁻¹) averaged across the four observers on the primary Y-axis and corresponding upper temporal frequency limits (Hz) on the secondary Y-axis. As can be observed, there is a general replication of the results observed by Morgan and Castet over the common range of spatial frequencies used in both experiments (0.08, 0.16, and 0.32 cycle/degree). Indeed, in terms of velocity, there is a proportional decrease in the estimated upper velocity limit for spatial frequencies ranging between 0.08 and 0.32 cycle/degree. Due to the proportional relationship between velocity and spatial frequency, when these data are expressed as temporal frequencies, a relatively constant upper temporal frequency limit of approximately 26 Hz is observed.

Upper velocity limits for spatial frequencies above 0.32 cycle/degree depict a different trend. As spatial frequency is increased beyond 0.32 cycle/degree, there is a non- proportional reduction in the upper velocity limit as indicated by the steeper slope in the velocity data. It should be noted that the bend in the velocity data is not apparent due to the large range of velocities offered by observer performance. When converted as a temporal frequency, the steeper slope for spatial frequencies above 0.32 cycle/degree results in a decrease in the upper temporal frequency limit as spatial

Figure 5

Figure 5 : Upper velocity limit (hollow symbols) and upper temporal frequency limit (solid symbols) as a function of spatial frequency (data averaged across four observers). Error bars represent two standard errors.

frequency is increased (17, 11, and 5 Hz for spatial frequencies of 0.64, 1.28, and 2.56 cycles/degree respectively). These results indicate that the visual system's temporal resolution for spatial frequencies of 0.64, 1.28, and 2.56 cycles/degree is approximately 35, 58, and 81% less than the temporal frequency limit of approximately 26 Hz observed for spatial frequencies in the 0.08 to 0.32 cycle/degree range.

Panels a, b, c, and d in figure 6, depict individual stereo performances across spatial frequencies as a function of temporal frequency (Hz). Figure 6 illustrates that the decreasing trend in upper temporal frequency depicted in figure 5 does not result from averaging of data across observers. Indeed, the gradual reduction in the upper temporal frequency limit with increasing spatial frequency is evident for all observers (albeit with individual differences).

An alternative way to conceptualise the decrease in the temporal resolution of the visual system for higher spatial frequencies, is to convert the upper temporal frequency limit as a phase angle displacement per frame (6.67 msec). For example, for spatial frequencies between 0.08 and 0.32 cycle/degree, an upper temporal frequency limit of 26 Hz represents a phase angle shift of approximately 62.9° per frame. However, as spatial frequencies increases beyond 0.32 cycle/degree, the maximum phase angle shift per frame at which depth can still be reliably reported

Figures 6 a and b

Figures 6 a and b : Relative disparity thresholds (simple gratings) as a function of temporal frequency

Figures 6 c and d

Figure 6 c and d : Relative disparity threshold (simple gratings) as a function of temporal frequency

decreases gradually to 40.2°, 13.2°, and 8.5° for spatial frequency gratings of 0.64, 1.28, and 2.56 cycles/degree respectively.

DISCUSSION

Results from Experiment I provide a replication and an extension of the results observed by Morgan and Castet to a wider range of spatial frequencies. Estimated upper temporal frequency limits in stereo performance illustrate a consistent pattern of results where the visual system illustrates a relatively constant temporal resolution of approximately 26 Hz for spatial frequencies at or below 0.32 cycle/degree. These data are consistent with those observed by Morgan and Castet for a similar range of spatial frequencies. However, our data also show that the temporal resolution of the visual system for encoding disparities in lateral motion becomes increasingly temporally low-passed as spatial frequency is increased beyond 0.32 cycle/degree. This is indicated by the observed reduction in the upper temporal frequency limit of stereo performance for spatial frequencies above 0.32 cycle/degree or, alternatively, a non-proportional reduction in the upper velocity limit as spatial frequency increases.

These results are consistent with those observed by Shoham et al. as well and Waugh and Hess, who demonstrated that the temporal frequency preference of the visual system for spatial frequencies within this range (i.e., 0.08 to 2.56 cycles/deg) becomes progressively more temporally low-pass with increasing spatial frequency. Prince et al., (1998) recently demonstrated that disparity information is processed by

distinct spatial frequency channels. Although our study did not seek to identify spatial frequency channels in stereopsis, our results are consistent with this view. However, results from Experiment I also suggest that these disparity channels are also likely to be tuned for temporal frequency.

The similarity in the reduction of the temporal resolution of the visual system for motion (Shoham et al., 1997; Waugh & Hess, 1994) and disparity processing in dynamic displays (Experiment I) with increasing spatial frequency, is also consistent with the notion that our results are due to a reduction in the temporal resolution of neurons which are responsible for the processing of both motion and disparity. As discussed earlier, there are a number of physiological and psychophysical studies which have examined the motion sensitivity of disparity detectors. These studies have demonstrated that neurons that are disparity selective also show a selectivity for the direction of motion (Bradley et al., 1995; Maunsell & Van Essen, 1983; Kwas et al, 1995). Given the physiological data suggesting a common neural substrate for the processing of motion and disparity, it would be expected that the ability of the visual system to maintain a disparity signal with increasing velocity would show a pattern of results which is consistent with data observed in motion (i.e., a decline in the upper temporal frequency limit with increasing spatial frequency).

It should be noted, however, that at the temporal frequency at which depth could no longer be reliably reported at threshold disparity, the direction of motion

could still be clearly perceived. This result is consistent with that observed by Morgan and Castet who also reported that observers could reliably report the direction of motion despite being unable to reliably report depth at threshold disparity. As suggested by Morgan and Castet, this may indicate that the upper velocity limit for depth and motion are not identical. Thus, although both disparity and motion may be processed jointly in the brain, the upper velocity limit for stereopsis is lower than that of motion.

EXPERIMENT II A

The goal of Experiment II is to assess whether or not interactions across spatial frequencies observed in motion and static stereo vision can also be observed in stereo performance using drifting compound gratings. Compound gratings were created with a one octave (adjacent compound gratings) and a two octave (non-adjacent compound gratings) separation between the high and the low spatial frequency component. As demonstrated by previous research (Wilson et al, 1991; Chang & Julesz, 1983; Eagle, 1997; 1998), it is expected that in adjacent compound gratings, spatial frequency components will interact in determining the upper velocity limit of stereopsis. The nature of this interaction, however may depend on the constituent spatial frequencies. For compound gratings comprising spatial frequencies at or below 0.32 cycle/degree, the upper velocity limit is expected to conform to that of the lower spatial frequency when tested alone. This is because, as demonstrated in Experiment I,

Morgan and Castet, and Shoham et al., these spatial frequencies are processed with similar temporal efficiency by the visual system.

For adjacent compound gratings with spatial frequencies processed with differing temporal resolutions, interactions could lead to two potential outcomes. From a strictly signal processing viewpoint, the visual system would gain in processing independently disparity signals at lower spatial frequencies. This would allow the visual system to rely on lower spatial frequencies which support higher velocities. Although this would represent the most efficient strategy, data from Cleary and Braddick, Chang and Julesz, and Eagle (1997; 1998) suggest that the visual system can not always take advantage of the best available signal with a one octave separation between spatial frequencies. Thus, for adjacent compound gratings with constituent spatial frequencies processed with different temporal resolutions, high spatial frequency constraints on the upper velocity limit of stereopsis could be observed.

For non-adjacent compound gratings, upper velocity limits are expected to conform to that of the lower spatial frequency. This is due to the general observation of a relative independence in the processing of spatial frequencies separated by two octaves. Thus, it is expected that the visual system will have independent access to disparity signals at lower spatial frequencies regardless of the particular combination of spatial frequencies.

METHOD

Stimuli and Procedure: The stimuli used in Experiment II A comprised compound gratings composed of the spatial frequencies used in Experiment I. Essentially, two types of compound gratings were generated. One set of compound gratings consisted of a high and low component separated by only one octave (0.16/0.32, 0.32/0.64, 0.64/1.28, and 1.28/2.56 cycles/degree) (adjacent compound gratings) and the other set comprised spatial frequencies separated by two octaves (0.16/0.64, 0.32/1.28, and 0.64/2.56 cycles/degree) (non-adjacent compound gratings). In order not to exceed available intensity values of the monitor, individual components of each compound grating were generated separately at 50% contrast. These components were then combined linearly as a single stimulus in the look up table. Except for differences in stimuli, the observers and the general methodology are the same as in Experiment I.

RESULTS

Static stereo thresholds:

Disparity thresholds for static adjacent compound gratings are presented in Figure 7. Also included in the figure are disparity thresholds for simple gratings (Experiment I). As can be observed, disparity thresholds for compound gratings are quite similar with those observed for simple gratings. These results are consistent with those observed by Smallman and MacLeod (1997) who demonstrated that

Figure 7

Figure 7 : Static disparity thresholds (arc min) for simple (hollow symbols) and compound (solid symbols) gratings as a function of spatial frequency. Error bars represent two standard errors.

disparity thresholds for compound gratings are equivalent to disparity thresholds for the lowest of the two components (see discussion Experiment II).

Dynamic stereo thresholds (adjacent compound gratings):

Data for each observer and for each compound grating were fitted with equation 1 (see panels a, b, c, and d in figure 8). Again, the function provided a good fit to the data with an average adjusted R squared value of .967 across conditions. Because compound gratings are constructed by generating each component separately and then combining them linearly in a look up table, the velocity expressed on the X-axis is common to both the high and the low spatial frequency component of the compound grating. Component temporal frequencies, however, differ markedly across spatial frequencies. Thus, it is more convenient to compare the upper velocity limit of compound and simple gratings as opposed to comparing their respective upper temporal frequencies. Also, by plotting the data as a function of velocity, we avoid a secondary X-axis which would be needed to represent the temporal frequency of the higher spatial frequency component.

Panels a, b, c, and d in figure 8 depict the stereo performance of observers SL, NC, LR, and JG as a function of velocity for each adjacent compound grating. In order to assess the presence of high spatial frequency constraints on the upper velocity limit of the lower component in adjacent compound gratings (hollow symbols), data from Experiment I depicting stereo performance for simple gratings are also included

Figures 8 a and b

Figures 8 a and b : Relative disparity thresholds for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity

Figures 8 c and d

Figures 8 c and d : Relative disparity thresholds for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity

(solid symbols). As can be observed, results for adjacent compound gratings with individual components at or below 0.32 cycles/degree show no reductions in the upper velocity limit when compared to the lowest component alone. These results are consistent with the idea that spatial frequencies processed with similar temporal resolutions will yield an equivalent performance when combined into adjacent compound gratings. However, these results contrast with stereo performance for adjacent compound gratings comprising a high and low spatial frequency with different upper temporal frequency limits. Indeed, for these compound gratings, stereo performance is intermediate to that of either spatial frequencies.

As in Experiment I, the fitted χ parameter of the function was used as an estimate of the upper velocity limit in stereo performance. Figure 9 depicts the average (four observers) upper velocity limit for each adjacent compound grating (black bars) along with the average upper velocity limit for simple gratings in Experiment I (white bars). As can be observed in figure 9, there is a clear tendency for the upper velocity limit of adjacent compound gratings comprising spatial frequency components below 0.32 cycle/degree to conform to the upper velocity limit of the lower component. As spatial frequencies in the compound gratings are increased however (above 0.32 cycle/degree), upper velocity limits tend to conform to a value which is intermediate to the low and high component. This indicates that although the ability of observers to reliably report depth at higher velocities is not

Figure 9

Figure 9 : Average (four observers) upper velocity limit for simple (black bars) and adjacent compound (white bars) gratings. Error bars represent two standard errors.

restricted to the poorer temporal resolution of the visual system at higher spatial frequencies, it is clear that the presence of the higher spatial frequency is constraining the upper velocity limit of the lower spatial frequency component for disparity judgements.

Dynamic stereo thresholds (non-adjacent compound gratings):

Upper velocity limits of stereo performance in non-adjacent compound gratings were unaffected by the presence of the higher spatial frequency component. Indeed, the higher and lower spatial frequency components appeared as a single stimulus over the lower range of velocities. As the velocity of the compound grating was increased, however, an independence of the individual components was observed, where the high and low components appeared to move transparently across each other with the higher component quickly becoming aliased and moving in the opposite direction. This decoupling of the motion signals across spatial scales, allowed observers to access the disparity signal of the lower component independently of the higher component. When observers were asked to report the depth (crossed or uncrossed) of the lower component, performance was unhindered by the presence of the higher spatial frequency and the upper velocity limit conformed to that of the lower component in Experiment I.

EXPERIMENT II B

In Experiment II A, the high and low component of the compound gratings were generated at 50% contrast and were then combined linearly in the look up table. This was done to ensure that the amplitude of the compound gratings did not exceed that of simple gratings. However, the components were combined in a relative phase angle that made the peaks and troughs of the individual component gratings coincide (peak added). There is a possibility that combining the constituent spatial frequencies in a peak added fashion may have created featural differences that could account for our data. This would suggest that the constraints observed in Experiment II A are not due to differences in the temporal resolution of the visual system to different spatial frequencies, but to an artifact owing to local differences in the compound stimuli when compared to simple gratings. Furthermore, a key assumption in Experiment II A is that the constituent spatial frequencies in the adjacent compound gratings are processed independently and that the reductions in the upper velocity limit observed in Experiment II A are due to interactions across these spatial frequencies. If this is the case, then the method of combining the two constituent spatial frequencies (i.e., peak added or peak subtracted) should have no influence on our data. To test this hypothesis, adjacent compound gratings were again created at 50% contrast but were combined in a peak subtracted fashion. However, testing was limited to the 0.32/0.64, 0.64/1.28, 1.28/2.56 cycles/degree adjacent compound gratings.

RESULTS

Results for the peak subtracted condition are presented for each observer in Panels a, b, c, and d in figure 10. Also included in the panels are the stereo performances for simple gratings (Experiment I). As can be observed, high spatial frequency constraints are still present using these stimuli. The data from Experiment II B were fitted with equation 1 (average adjusted R squared value of 0.974) and average estimated upper velocity limits are presented in Figure 11 (grey bars) along with relevant data from Experiment I and II A (black and white bars respectively). As can be observed in figure 11, it is clear that stereo performance in compound gratings is not dependent on featural differences resulting from the higher contrast generated by the peak added compound gratings. Indeed, high spatial frequency constraints are readily observable whether we consider the peak added or peak subtracted compound gratings. From these results, it can be concluded that both forms of compound gratings yield a similar pattern of results and that the amplitudes of the individual spatial frequency components, not the overall contrast, is the determining factor in the observed stereo performance.

The fact that disparity processing in compound gratings are not contingent on component phase has also been observed in other experiments (Rohaly & Wilson, 1993; Schor, Heckman, & Tyler, 1989). For example, Rohaly and Wilson have assessed the effect of component phase on disparity averaging for gratings of different

Figure 10 a and b

Figures 10 a and b : Relative disparity thresholds for simple (solid symbols) and compound peak subtracted (hollow symbols) gratings as a function of velocity

Figure 10 c and d

Figures 10 c and d : Relative disparity thresholds for simple (solid symbols) and compound peak subtracted (hollow symbols) gratings as a function of velocity

Figure 11

Figure 11 : Average upper velocity limit (four observers) for simple (black bars), peak added compound (white bars), and peak subtracted compound (grey bars) gratings. Error bars represent two standard errors.

spatial frequencies. Their results demonstrated no systematic differences on the averaged depth plane resulting from the relative phase of the component gratings. Schor et al. also examined the effect of relative component phase of compound gratings on the fusion limit. Their results revealed no differences across relative phase conditions on the fusion limit of stereopsis. Our results along with those of Rohaly and Wilson and Schor et al. are more consistent with the notion that disparities across spatial frequencies are processed independently and that stereo performance in compound gratings does not result from featural differences resulting from the relative phase of the component spatial frequencies.

DISCUSSION

Results from Experiment II demonstrated that the upper velocity limit of stereopsis for drifting compound gratings is subject to high spatial frequencies constraints. These constraints however are limited to adjacent compound gratings with low and high spatial frequencies which are resolved with differing temporal efficiency by the visual system. These results are in agreement with a number of studies in motion and stereopsis that have demonstrated that for stimuli comprising adjacent spatial frequencies, the visual system may not be able to tease individual components apart (Chang & Julesz, 1983; Eagle, 1998; Wilson et al., 1991; Rohaly & Wilson, 1993). As a result, when computing the depth of a drifting compound

grating, the visual system must contend with a clasped drifting disparity signal arising from spatial frequencies which are processed with differing temporal resolutions.

The independence of individual constituent spatial frequencies in non-adjacent gratings is also consistent with results observed in stereopsis and motion where, for example, it has been shown that noise presented at higher spatial frequencies will only affect stereopsis at lower spatial frequencies if the noise and disparity signals are in adjacent spatial frequencies (Julesz & Miller, 1975; Prince et al., 1998). This has also been observed in motion where higher spatial frequency constraints on d_{max} have been shown to disappear if the RDK's middle spatial frequencies are filtered out (Chang & Julesz, 1983; Eagle, 1998).

More recently, Qian, et al., (1994) have shown that transparent motion can not be induced with superimposed gratings of identical spatial frequency. With such stimuli, observers typically perceive either periodic oscillation of the whole pattern or flicker when the temporal frequency is high. The flicker in the gratings occurs without any net motion in either direction suggesting that the individual motion of the two components cancel each other out (motion-opponency effect). For superimposed gratings separated by two octaves however, transparent motion was always perceived. Our results suggest that transparent motion across components in non-adjacent gratings does not occur at the cost of disparity. In fact, despite having two gratings

moving transparently across each other, the disparity signal at the lower spatial frequency could still be accessed.

Overall, results from Experiment II for adjacent compound gratings illustrate a lawful pattern of results based on the visual system's temporal resolution at individual spatial frequencies. As the literature suggests, in processing motion and disparity, the visual system is unable to tease apart adjacent components. A natural extension of these results would suggest that disparity and motion signals for stimuli comprising adjacent spatial scales are clasped or averaged into a single signal. Under such circumstances, the ability of the visual system to maintain a disparity signal at high velocities may become contingent on the pooled temporal responses across spatial frequencies. In this case, disparity signals arising from a pooled response of spatial frequencies which are resolved with a similar temporal resolution, would show an equivalent response when combined into a single stimulus. However, disparity signals arising from the pooled responses of spatial frequencies resolved at differing temporal resolutions would be constrained by the temporally less efficient higher spatial frequencies³.

3

We also replicated the first and second experiments with the use of a septum (observer SL and NC) so as to ensure that the constraints observed in experiment II were not due to the low luminance of the display resulting from the use of polarized sheets (peak and trough luminance of 9.25 and 0.34 cd/m²). Results of these control experiments demonstrated that high spatial frequency constraints were also present at higher luminance levels (peak and trough luminance of 88.9 and 3.2 cd/m²).

The question arises as to why motion signals arising from adjacent spatial frequencies are not processed independently. From an ecological perspective, such a pooling of responses would prove useful in binding various motions signals to their appropriate objects in the visual environment. Indeed, a stimulus with components in (1) adjacent spatial frequencies, (2) on a common depth plane, and (3) moving at a common velocity may indicate to the visual system that these signals are arising from a common object. It is possible that in binding appropriate motion and disparity signals together, the visual system preserves a coherent percept of the visual environment. It is interesting to note that the literature in motion and disparity suggest that for compound gratings, a difference in disparity and/or spatial frequency would cause motion signals to unfasten. This is likely to occur because objects in the visual environment rarely simultaneously occupy two widely distinct depths planes or have spatial scales at both ends of the spectrum without any middle frequencies.

Although our results do not provide insight into the nature or presence of high spatial frequency constraints in static displays, they do provide empirical evidence that such constraints are being exercised in dynamic displays. As argued above, these constraints likely occur as a result of a binding of motion and disparity signals arising from spatial frequencies which are resolved with differing temporal resolutions by the visual system. However, whether or not such constraints are observed in stereopsis depends largely on the relationship entertained between key elements of a display.

Indeed, high spatial frequency constraints are more likely to be observed across adjacent spatial frequencies and on a shared depth plane.

Signal to Noise Ratio Due to Display Aliasing

We wondered as to an alternative cause for the constraints observed in adjacent compound gratings. One possibility is that the signal arising from the higher spatial frequency component may have introduced noise at very high velocities due to display aliasing or a loss in signal amplitude with increasing temporal frequency (i.e., a display artifact). Disparity, like all visual signals, must be derived in the presence of noise which is inherent to the visual system. If the disparity in compound grating is computed by considering jointly the disparity signals at the high and low spatial frequencies, it is reasonable to suggest that the disparity signal at higher spatial frequencies becomes increasingly noisy as velocity increases. This could have had an additive effect on the internal noise of the visual system. Thus, the reduction in the upper velocity limit may simply represent the point at which the overall disparity signal for the compound grating is, in proportion, smaller than the sum of the internal and the noise generated by the display.

A number of results from Experiment I seem to contradict this explanation. Aside from the fact that observers did not report perceiving motion in the direction opposite to the intended direction, from a temporal frequency perspective, a 0.16/0.32 cycle/degree compound grating is identical to a 1.28/2.56 cycles/degree compound

grating. Hence, if increased display noise from the higher spatial frequency was the explanation for the results observed in Experiment II, constraints over the 0.08 to 0.32 cycle/degree range should also have been observed. This, of course, is not the case. Indeed, our results demonstrate that combining spatial frequencies within this range into a compound grating has no effect on the upper velocity limit of the lower component.

Another result that seems to contradict this explanation is that, beyond 0.32 cycle/degree, the upper temporal frequency limit decreases in a fashion that is no longer proportional to the spatial frequency of the stimulus. From a display perspective, these results imply that the demands being exacted on the display to generate an upper velocity limit decreases with increasing spatial frequency. In other words, the "noise" generated by the display decreases as stimulus spatial frequency increases. This runs counter to our results which indicate that high spatial frequency constraints are greater at higher spatial frequencies. Finally, we assessed the temporal response of the monitor by placing a photo sensitive cell over a slitted opaque square piece of plastic while the gratings were drifted at 10, 20, 35, and 40 Hz. The signals were then digitised and sampled at 5000 samples per second. Comparisons in the amplitude of the signal across temporal frequencies showed no appreciable loss of amplitude with increasing temporal frequency.

Size-Disparity Correlation

Results from Experiment II seem to be in strict contradiction with the low spatial frequency constraint hypothesis whereby stereo matches at lower spatial frequencies are derived independently and then serve to constrain stereo matches at higher spatial frequencies (Marr & Poggio, 1979; Wilson et al. 1991). For laterally moving disparity information, such a scheme would suggest that disparity information at lower spatial frequencies would be available prior to and independently from disparity information at higher spatial frequencies. Given this, an efficient response to disparity signals moving at high velocities would imply that the visual system would rely on signals arising from lower spatial frequencies as they support proportionately higher velocities. In fact, the results observed in Experiment II suggest, in a very consistent way, that interactions across spatial frequencies (i.e., adjacent compound gratings) illustrate mainly a constraint of lower spatial frequencies by higher spatial frequencies.

Unfortunately, empirical evidence for a high spatial frequency constraints on stereopsis is rather weak. Indeed, the bulk of the argument for such a scheme in disparity coding has relied mainly on the failure to confirm the low spatial frequency constraint hypothesis in certain experimental procedures. For example, Smallman and MacLeod (1997) have argued that if stereo matches at lower spatial frequencies are used to drive stereo matches at higher spatial frequencies, then disparity thresholds for

a compound grating should be better than the disparity threshold of the lower component alone. Overall, experiments that have examined this issue have provided mixed results with disparity thresholds for compound gratings which are sometimes marginally better, sometimes worse, but most often the same as in simple gratings (note that our results show that disparity thresholds for compound gratings-Experiment II-were no different than those of the lower component alone-Experiment I-). Although such results may be interpreted as detrimental to the low to high constraint hypothesis, they do not necessarily imply a high to low constraint.

Other results, such as those observed in Smallman's disambiguation experiment, demonstrate that high spatial frequencies can be used to disambiguate lower spatial frequencies. Admittedly, however, this represents a special case as the lower spatial frequency used in the experiment was completely ambiguous and thus provided no disparity signal. With no disparity signal in the lower spatial frequency, it is not surprising that the ambiguous low spatial frequency conformed to the unambiguous disparity signal of the higher spatial frequency.

There is one psychophysical study that seemed to have little problem in demonstrating high spatial frequency constraints in stereopsis. In a previous section, a study conducted by Rohaly and Wilson was discussed where they asked observers to view a dichoptic pair of spatial frequency gratings superimposed on a second pair with a different spatial frequency. One of the two pairs (either the low or the high spatial

frequency pair) was presented at zero disparity while the disparity of the other pair was varied. Observers were then asked to compare the perceived depth of these superimposed gratings with that of an identical reference display in which the disparity of both the low and high spatial frequencies were equal. In other studies (Kaufman, Bacon, & Barosso, 1973; Foley, 1976; Parker & Yang, 1989) it had been shown that if the difference in disparity between superimposed dichoptic pairs of gratings is small, there is a tendency for the visual system to pool the two disparity signals and yield a common intermediate depth plane (a phenomenon called depth pooling or disparity averaging).

Their results demonstrated that the high spatial frequency component seemed to dominate the averaging process. In fact, observers consistently reported either an arithmetic average disparity or a depth plane which was closer to the actual disparity of the higher spatial frequency component. As noted by Rohaly and Wilson, these results seem to be in strict contradiction with those of Wilson et al. who demonstrated that lower spatial frequencies constrained the fusion limit of higher spatial frequencies but that the reverse was never observed.

To explain these discrepancies, Rohaly and Wilson argued that the change in the direction of spatial frequency interactions across experiments could be accounted for by the fact that they used mainly small disparities (i.e., 112 arc sec) and that Wilson et al. used large disparities (approximately 14.5 arc min). The suggestion is

that using large disparities tend to favour low spatial frequencies, whereas using small disparities tend to favour high spatial frequencies. This analysis is in agreement with the general observation of a size-disparity correlation in psychophysical studies which has lead other experimenters to suggest that large disparities are processed by low spatial frequency channels and that small disparities are processed by high spatial frequency channels.

If there is a link between the processing of disparity and spatial frequency, both the stereo threshold and the upper disparity limit of stereopsis should be related to the spatial frequency of the stimulus. As discussed earlier, there is good evidence of these trends in the data provided by Legge and Gu for disparity thresholds and those provided by Smallman and MacLeod (1994) for the upper disparity limit of stereopsis. Schor, Wood, and Ogawa (1984) have also investigated this question using small patches of difference of Gaussian (DOG) luminance profiles with a centre spatial frequency ranging between .075 cycle/degree and 19.2 cycles/degree. Their results, much like those observed by Legge and Gu and Smallman and Macleod (1994) show that both disparity thresholds and the upper disparity limits of stereopsis decrease with increasing spatial frequency. The decrease for disparity thresholds, however, increased slightly more rapidly than the upper disparity limit. This indicates that the range of disparities over which stereopsis is possible decreases as spatial frequency increases.

Another prediction made by a size-disparity correlation is that stereo thresholds at higher spatial frequencies should fall off at a faster rate as we move away from the fixation plane. This is because disparities get larger as we move away from the fixation plane. Given this, a size-disparity correlation would imply that lower spatial frequencies are more efficient at processing disparities as the target is moved away from fixation. Smallman and MacLeod (1997) examined this particular issue in a depth discrimination experiment. Depth discrimination thresholds differ from the usual disparity thresholds in that they require that observers make judgements about the relative depth of two stimuli (a reference and test stimulus) when both presented away from the fixation plane. Usually, the reference stimulus is presented at a fixed disparity (standing disparity), while the disparity of the test stimulus is varied. Thresholds in these displays are defined as the test disparity needed to perceive the test and reference stimuli as appearing on separate depth planes. Smallman and MacLeod (1997) presented pairs of filtered texture patterns with centre spatial frequencies varying between 1 and 11 cycles/degree. The reference stimulus was presented at standing disparities ranging between 5 and 30 arc min in steps of 5 arc min. Depth discrimination thresholds were then assessed for test patterns with spatial frequency content identical to the reference stimulus.

Results from their experiment showed that for small standing disparities (i.e., close to the fixation plane) depth discrimination thresholds were similar across spatial

frequency, with a slight advantage for higher spatial frequencies. As the stimuli were moved away from fixation (i.e., larger standing disparities), however, depth discrimination thresholds increased dramatically for higher spatial frequencies. For example, for a test stimulus of 11 cycles/degree, an increase in the standing disparity from 4 to 12 arc min required a 2000% increase in disparity of the test stimulus to generate depth discrimination. In comparison, for an equivalent difference in standing disparity, a stimulus with centre spatial frequency of 2 cycles/degree required only a 30% increase in disparity. What these results illustrate is that high spatial frequencies are not efficient in processing large disparities. These results are consistent with a size-disparity correlation.

Implementation of a size-disparity correlation in the processing of disparity information could have functional consequences on the upper velocity limit of stereopsis. In keeping with the general approach suggested by Marr and Poggio (1979) and Prince et al. by which disparity is processed via spatial frequency channels, imagine spatial frequency channels tuned to spatial frequency and disparity. If small disparities are processed by channels tuned to higher spatial frequencies, then a compound grating comprising a high and low spatial frequency presented at a small disparity would generate a greater response from channels tuned to the higher spatial frequency. However, channels tuned to lower spatial frequencies would be contending with a disparity that is further away from their preferred disparity, thus

yielding a moderate or weak response. As the compound grating is put into motion, the visual system must process signals arising from adjacent spatial frequencies, on a common depth plane, and moving at a common velocity. This induces the visual system to process information stemming from these different spatial frequencies as if they were arising from a common object. In turn, this leads to a compromise between motion signals whereby overall performance is mitigated by the presence of a strong signal arising from the more sluggish/higher spatial frequency channels. If larger disparities are introduced in the display, channels tuned to the lower spatial frequencies will present a stronger signal as the disparity is closer to their preferred disparity. For channels tuned to the higher spatial frequencies however, larger disparities may fall outside of the range of preferred disparities, yielding a weaker response. Under such circumstances, we would expect little influence of these less temporally efficient spatial frequency channels in the binding of motion signals. That is, a release from high spatial frequency constraints could be expected at larger disparities.

Taken together, results obtained by Rohaly and Wilson and Wilson et al., Legge and Gu, Smallman and MacLeod (1994); Smallman and MacLeod (1997) suggest that the high spatial frequency constraints observed in Experiment II may have resulted from the relatively small disparities used (ranging between 1.4 and 5 arc min over the range of spatial frequencies where constraints were observed). If indeed,

smaller disparities are processed by high spatial frequency channels, it is possible that in adjacent compound gratings, the contribution of the high and low spatial frequencies may have been slightly in favour of higher spatial frequencies. We wondered if it was possible to alter the relative contribution of signals arising from different spatial scales by increasing the overall disparity of the stimuli. In order to test this hypothesis a third experiment was conducted with the stimuli used in Experiments I and II, but at a greater disparity.

EXPERIMENT III

METHOD

Stimuli and Procedure: The stimuli and display equipment used in Experiment III were identical to those used in Experiments I and II. For compound gratings however, testing was limited to adjacent components. A static disparity threshold was estimated at which observers could accurately report the depth of the static grating at 100% correct response rate. This was done by setting the disparity to threshold values observed in Experiments I and II and gradually increasing the disparity of the display in steps of one pixel until observers performed at a 100% correct response rate on 30 consecutive trials. This procedure was repeated three times and the average disparity observed over the three sessions was used as an estimate of the disparity yielding a performance just at 100%. Once disparities yielding a 100% correct response rate using static displays were obtained, observers completed 60 trials at each velocity in

separate blocks of trials (again with increasing velocities). Responses were recorded as being correct or wrong on each trial, and the appropriate proportion of correct responses was derived at each velocity. Thus, instead of defining the upper velocity limit of stereo performance as an increase in the relative disparity threshold, the proportion of correct responses was assessed as a function of velocity. Thus, for this experiment upper velocity limits are defined as the point at which stereo performance declined as a percentage of correct responses.

RESULTS

Static Disparity Thresholds:

Figure 12 presents the estimated disparity yielding a 100% correct response rate for each observer and at each spatial frequency. Also included in the figure are disparity thresholds from Experiment I. Generally, asking observers to perform at a 100% correct response rate, generated a doubling of the disparity threshold observed in Experiment I. As expected, the disparity thresholds observed in Experiment III also conform to an inversely proportionate relationship with spatial frequency.

Simple Grating :

P correct responses for simple gratings as a function of velocity are presented in panels a, b, c, and d in figure 13 for the four observers and each spatial frequency. For comparative purposes, relevant data from Experiment I are also presented. The data from Experiment III were fitted with a modified version of equation 1:

Figure 12

Figure 12 : Static disparity thresholds for 75% (hollow symbols) and 100% (solid symbols) correct responses. Error bars represent two standard errors.

Figures 13 a and b

Figures 13 a and b : P correct (hollow symbols) (Exp III) and relative disparity thresholds (solid symbols) (Exp I) as a function of velocity

Figures 13 c and d

Figures 13 c and d : P correct (hollow symbols) (Exp III) and relative disparity thresholds (solid symbols) (Exp I) as a function of velocity

$$PC = \beta * [1 - \exp^{-\alpha(v-\chi)}] \quad \text{Equation 2.}$$

Where PC = proportion of correct responses, v is velocity in deg sec⁻¹, and α , χ , and β are constants representing the rate of growth, point at which performance decline, and asymptotic performance respectively. Again, the function was shown to provide a good fit to the data with an average adjusted R squared value of .975.

As figure 13 illustrates, deterioration in stereo performance with increasing velocity occurs at approximately the same velocity regardless of whether we consider relative threshold or proportion of correct responses as a criterion. This, however, is not the critical analysis. As in Experiment I, the important point is that, the temporal frequency at which depth can no longer be reliably reported decreases as spatial frequency increases. To assess this, observer performance was replotted against temporal frequency. Panels a, b, c, and d in figure 14 depict p correct performance as a function of temporal frequency. As in Experiment I, a gradual decline in the upper temporal frequency with increasing spatial frequency is observed for all observers.

Adjacent Compound Gratings :

Panels a, b, c, and d in figure 15 depict the proportion of correct responses of each observer for simple (solid symbols) and compound (hollow symbols) gratings.

Figures 14 a and b

Figures 14 a and b : P correct as a function of temporal frequency

Figures 14 c and d

Figures 14 c and d: P correct as a function of temporal frequency

Figures 15 a and b

Figures 15 a and b: P correct (primary Y-axis) and relative disparity threshold (secondary Y-axis) for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity.

Figures 15 c and d

Figures 15 c and d: P correct (primary Y-axis) and relative disparity threshold (secondary Y-axis) for simple (solid symbols) and compound (hollow symbols) gratings as a function of velocity.

Again, relevant data from Experiments I and II (peak added) are included for comparative purposes. As in Experiment II, compound gratings with both the high and low spatial frequency component at or below 0.32 cycle/degree yield no high spatial frequency constraints on the upper velocity limit of the lower spatial frequency component. Thus, within this range of spatial frequencies, results from Experiment III represent a replication of the results observed in Experiment II. Interestingly, however, beyond 0.32 cycle/degree, higher spatial frequency constraints observed in Experiment II at smaller disparities are, for all intents and purposes, largely eliminated. In fact, except for observers NC and JG with a 1.28/2.56 cycles/degree compound grating, there is little difference in the performance afforded by adjacent compound gratings and that of the lower component alone.

For observer NC and JG at the 1.28/2.56 cycles/degree compound grating condition, although parameter estimates of the fitted functions across the compound and simple grating conditions show that performance across these conditions do not differ, the fact that potential constraints in this condition were observed in 2 out of 4 observers leads one to be tentative about concluding that constraints were fully eliminated in this condition. Nonetheless, it is important to point out that, although constraints may still be present in this condition, they are greatly reduced when compared to the constraints observed in Experiment II for the same condition.

DISCUSSION

For simple gratings, results from Experiment III show that reductions in the upper temporal frequency limit for higher spatial frequencies are still present when larger disparities are introduced in the displays. These results are not entirely unexpected. Indeed, although disparities in Experiment III were larger (approximately twice that of Experiments I and II), the performance required of observers was higher as well (i.e. 100% correct compared to 75% correct).

The more interesting result, however, is that increasing disparity largely eliminated higher spatial frequency constraints on the upper velocity limit of the lower component in adjacent compound gratings. Although such constraints have been observed in motion (Cleary & Braddick, 1990) and in stereopsis (Experiment II), to our knowledge these results represent the first demonstration that lower spatial frequencies can be relieved of these constraints by simply presenting the stimulus further away from the fixation plane.

The fact that performance in adjacent compound gratings at larger disparities conforms to the upper velocity limit of the lower component strongly suggests that performance in these displays is mainly driven by the lower spatial frequencies. Given the observed domination of lower spatial frequencies, results from Experiment III provide empirical evidence for the analysis made by Rohaly and Wilson who suggested that the direction of the interaction across spatial frequencies can be altered

by introducing either small (high spatial frequency constraint) or large (low spatial frequency constraint) disparities.

Adjacent compound gratings used in Experiments II and III have laid out conditions which are favourable to the observation of spatial frequency interactions (i.e., adjacent spatial frequencies and a common depth plane). In fact, aside from being further away from the fixation plane, the compound gratings used in these experiments were perceptually identical. Yet in Experiment III, high spatial frequency constraints were largely eliminated. These results suggest that disparity information plays a critical role in determining which spatial frequency will dominate in the pooling of motion and disparity signals.

Previous experiments (Bradley et al, 1995; Kwas et al. 1995) have shown that combining motion and disparity information in experimental displays provides the visual system with critical information about how to process visual stimuli. For example, it has been shown that disparity information is used by the visual system to determine which motion signals should be bound, and which should be treated separately. Our results provide empirical evidence to the notion that disparity not only serves to disentangle motion signals across depth planes, but can also serve to alter the nature of the interaction across spatial frequencies when binding motion signals.

GENERAL DISCUSSION

In Experiment I of this thesis, it was demonstrated that at very low spatial frequencies (0.08 - 0.32 cycle/degree), the upper temporal frequency limit of stereo vision is relatively constant and that this upper temporal frequency limit decreased as spatial frequency increased beyond this range. This replicated and extended the earlier work of Morgan and Castet who also observed a constant upper temporal frequency limit of 30 to 40 Hz at low spatial frequencies (0.04 to 0.32 cycle/degree).

A review of the literature suggests that there are strong interactions between spatial frequencies within a one octave range, whereas beyond the one octave range, spatial frequency information tends to be processed independently. One hypothesis that emerged from this, is that for compound gratings comprising two different spatial frequencies, the upper velocity limit of stereo vision should exhibit interactions across spatial frequencies when they are within a one octave range but should be free from interactions beyond this range. We hypothesised that the upper velocity limit of compound gratings comprising spatial frequencies with a one octave separation could show interactions such as a reduced upper velocity limit due to the presence of higher spatial frequencies. However, for compound gratings with a two octave separation little interaction would be expected and the upper velocity limit should conform to the spatial frequency yielding the best temporal response (i.e., the lower spatial frequency component).

Results from Experiment II demonstrated that for gratings comprising spatial frequencies with a one octave separation with different temporal resolutions, the upper velocity limit was below that offered by the lower spatial frequency. For compound gratings with spatial frequencies separated by one octave but with similar temporal resolutions, the upper velocity limit was equal to the lower/fastest spatial frequency. However, for compound gratings with a two octave separation, the upper velocity limit was always similar to the lower component. These results were taken as empirical evidence of high spatial frequency constraints on the processing of disparity in lateral motion for compound gratings with a one octave separation.

Experiment III examined the hypothesis that spatial frequency interactions could be modulated by the size of the disparity in the displays. Rohaly and Wilson suggested that at small disparities, high spatial frequencies may dominate processing, whereas at large disparities processing may be dominated by low spatial frequencies. This hypothesis rests on the assumption that the visual system favours high spatial frequencies for small disparities whereas it favours low spatial frequencies at larger disparities (i.e., size-disparity correlation). Hence, the relative contribution of spatial frequencies in determining the depth of a compound grating may be different depending on the size of the disparity. Accordingly, at small disparities, one would expect that the temporal processing of a compound grating would be dominated by the slower/higher spatial frequencies, whereas at larger disparities, we would expect the

temporal processing of the compound grating to be dominated by the faster/lower spatial frequencies. In order to test this hypothesis, a third experiment where larger disparities were introduced in the displays used in Experiment II was conducted. Consistent with Rohaly and Wilson's analysis, the results of Experiment III demonstrate that high spatial frequency constraints were eliminated. That is, regardless of the combination of constituent spatial frequencies, the upper velocity limit of the compound grating was always determined by the spatial frequency channel with the better temporal resolution.

Ecological Perspective

In this section, we will discuss the ecological relevance of the results observed in this thesis. Specifically, a link is made between optic flow processing and motion signal averaging across spatial frequencies. It is suggested that averaging of motion signals across spatial frequencies at various disparities may contribute to disentangling the translation and rotation component of the optic flow and thus provide critical information to observers when navigating in the visual world.

The contemporary view of stereo and motion processing is that motion and stereo are multiplexed within single units in the visual system. This view is supported by the general observation that, in area MT of the macaque monkey, neurons selective to direction of motion are also selective for disparity. The functional consequences of cooperation between motion and disparity in area MT has been demonstrated by

Bradley et al. who found that motion transparency was dependent on disparity information. It was previously demonstrated that the response of MT neurons to a random-dot pattern moving in its preferred direction was greatly reduced by the presentation of a second transparent dot pattern moving in the opposite direction. The important contribution of Bradley et al. is that these reductions in the response occur only when the disparity between preferred and anti-preferred patterns is within a certain range of one another. When the stimuli are clearly separated in depth, no reductions in the neurons response are observed. Based on these results, Lappe (1996) recently argued that the interaction between disparity and motion in area MT serves mainly to separate motion signals that originate from different depth planes. However, when motion signals are presented on similar depth planes, an averaging of motion signals occurs. The importance of disparity in processing of motion has also been demonstrated by Kwas et al. who showed that pattern motion in plaid stimuli could be observed only when plaid components are presented on a common depth plane.

As argued by Lappe, spatial averaging of motion signals across similar depth planes may play an important role in reducing the computational demands in processing optic flows resulting from an observer's motion in the visual environment. In computing accurately the direction of heading, the visual system must be able to

differentiate motion signals in the optic flow arising from translation direction from those arising from an observer's rotation.

In linear motion where observers look directly ahead while traveling, the translation component of the optic flow is reduced to the radial structure of motion signals arising from an observer's self motion in the environment. However, observers often navigate in the environment while examining objects that are not directly in line with the direction of heading. This results in motion signals that are not necessarily related to the direction of heading. Hence, the optic flow often comprises a second source of motion signals related to the observer's eyes, head, or body rotation (i.e., rotation component). Yet, despite the presence of these two motion components, observers are able to perceive accurately their direction of motion. As argued by Van den Berg and Brenner (1994), it follows that observers must be able to distinguish between motion signals arising from the translation and rotation components of the optic flow.

Lappe noted that, in the translation component, a natural parsing of motion signals and depth information occurs. That is, static objects close to each other in depth tend to move at similar velocities on the retinae, whereas static objects separated in depth tend to move at different velocities. From this, he argued that a spatial averaging of motion signals within a defined range of disparities could significantly reduce the computational demands of processing the translation

component of the optic flow. This would be achieved by claspings a number of motion signals on similar depth planes into a single velocity estimate.

However, there is a distinction to be made between the account proposed by Lappe and the motion signal averaging observed in this thesis. Lappe's account suggests an averaging of motion signals across adjacent depth planes. Spatial frequency components in our compound gratings shared a depth plane and averaging occurred across spatial frequencies. This points to the complex nature of interactions between stereo and motion. Indeed, although motion signals can be processed independently if there is enough distance between their respective depth planes (Bradley et al, 1995), they can also be processed independently on a common depth plane if there is enough distance between their spatial frequencies (Qian, Anderson, & Adelson, 1994; Experiment II of the current thesis). It is possible that when the conditions for motion signal averaging are met in depth (similar depth planes) and in spatial frequency (similar spatial frequencies), motion averaging is either reflected as an average velocity across depth planes or an intermediate upper velocity limit across spatial frequencies.

The question arises as to why the visual system would be designed in a manner in which motion signal averaging would show a compromise across spatial frequencies at small disparities whereas no such compromise is observed at larger disparities. One possibility is that, in processing the optic flow, motion averaging

across depth planes facilitates the processing of the translation component whereas motion averaging across spatial frequencies facilitates the processing of the rotation component. As discussed above, the rotation component of the optic flow results from navigating in the environment while examining surrounding objects. Looking at an object in the environment while moving around, often involves keeping its image on the fovea by rotating our eyes, head, and/or body. Keeping an object on the fovea reduces its motion on the retinae. However, visually tracking an object in this way, also causes objects further away from the fixation plane to move at higher velocities (motion parallax). In conjunction with this increase in velocity, disparities also become larger with increasing distance from fixation. From this, we can readily see that, in addition to the natural parsing of depth and motion in the translation component, there is also a velocity-disparity correlation in the rotation component.

As discussed earlier, there is an important difference in the nature of motion signals arising from the translation component and the rotation component. Strictly speaking, while the translation component refers only to the radial structure of motion signals resulting from observers moving towards a target location, the rotation component also includes strong lateral motion signals. Because of this distinction, it is possible to create a situation where the optic flow is mostly rotational in nature. For example, figure 16 illustrates a scenario where an observer would process an optic flow in which the translation component is minimal. As can be observed, the

Figure 16

Figure 16 : Observers often fixate an object as they navigate in the visual environment. As a result, objects near the fixation plane have small disparities and move at low velocities, whereas objects further away from the fixation plane have large disparities and move at high velocities

correlation between velocity and disparity is mainly due to lateral motion. Because the observer is fixating an object while he/she is moving, objects laying near the fixation point (small disparities) show little motion on the retinae, whereas objects further away from the fixation plane move at greater velocities. Accordingly, it is sensible that large disparities would be processed by channels tuned to low spatial frequencies as they support higher velocities. Conversely, objects laying near the fixation plane would be better processed by channels tuned to higher spatial frequencies where maximum spatial resolution is afforded and where the need for temporal resolution can be minimized by visual tracking. It is possible that by using disparity information as a means to weigh appropriately the contribution of low and high spatial frequencies in averaging motion signals, the visual system may obtain a robust estimate of the rotation component of the optic flow which can then be extracted from the optic flow to compute direction of heading.

Suggested Experiments

Perhaps the most novel result of the current thesis is the demonstration that disparity information can influence how motion signals interact across spatial scales. From this, a number of predictions can be made about how previously uncovered phenomena in stereopsis could be further explored by manipulating disparity.

Inter Ocular Contrast Differences

In a study, Cormack et al. (1996) showed that in broadband stimuli, stereo thresholds were unaffected by inter-ocular differences in contrasts. This was an unexpected result as previous research (Legge & Gu, 1989; Schor & Heckman, 1989; Halpern & Blake, 1988) demonstrated that such manipulations were extremely detrimental to stereo thresholds at lower spatial frequencies. As argued by Cormack et al. (1997), it is likely that the presence of higher spatial frequencies in broadband stimuli may have allowed the visual system to overcome the difference in interocular contrast that had proved so detrimental to stereopsis at lower spatial frequencies. Cormack et al. (1997) examined this hypothesis using spatial frequency gratings of various spatial frequencies and demonstrated that inter-ocular differences in contrast only had an effect on stereo thresholds at lower spatial frequencies. Based on these results, Cormack et al. (1997) attributed the absence of an inter-ocular contrast effect in broadband stimuli to the presence of higher spatial frequencies in the display. Unfortunately, the study of Cormack et al. (1996) using broadband stimuli, only examined stereo thresholds and were by definition limited to small disparities. An experiment examining the percentage of correct responses for front and behind judgements could be used to examine the effect of inter-ocular contrast differences on performance in broadband stimuli at large disparities. If by introducing larger disparities we generate disproportionate responses across spatial scales that are

favourable to lower spatial frequencies, we would expect that, even with broadband stimuli, performance would decline with increasing differences in contrast across the two eyes.

Disparity Averaging

Disparity averaging is the phenomenon by which two stimuli with different but similar disparities are perceived as a single stimulus lying on an intermediate depth plane. As discussed above, Rohaly and Wilson recently demonstrated that, at small disparities, the intermediate depth generated by the two stimuli seem to lie on a depth plane that is closer to the actual disparity of the higher spatial frequency. Rohaly and Wilson alluded to the possibility that using larger disparities may change the direction of the interaction across spatial scales (an assertion for which empirical evidence is provided in Experiment III). However, they limited the influence of the lower component to determining the fusion range of higher spatial frequencies. Our results suggest that the influence of larger disparities could also be potentially extended to influencing the intermediate depth plane in disparity averaging. That is, using larger disparities may yield an average depth plane closer to the actual disparity of the lower spatial frequency component.

Motion Transparency

As discussed above, Bradley et al. recently demonstrated that motion-opponent inhibition did not occur if the preferred and anti-preferred stimuli were presented on

different depth planes. A natural extension of the results of this thesis to studies in transparent motion would suggest that inducing transparent motion between high spatial frequency gratings would require a larger disparity difference between the gratings as they are moved further away from the fixation plane. That is, although relatively small differences in disparity between the gratings would suffice to induce transparent motion near the fixation plane, moving the pair of stimuli away from the fixation plane would require larger disparity differences. This is because higher spatial frequencies are less sensitive to differences in disparity signals at large standing disparities.

That such results would bear out is supported by the study conducted by Smallman and MacLeod (1997). To recap, Smallman and MacLeod (1997) examined depth discrimination thresholds at various distances from the fixation plane (standing disparities). Again, depth discrimination thresholds require that observers make judgements on the relative depth of two stimuli (a reference and test stimulus) when both are presented away from the fixation plane. Thresholds in these displays are defined as the test disparity needed to perceive the test and reference stimuli as appearing on separate depth planes. Their results revealed that as the stimuli were moved away from the fixation plane, depth discrimination thresholds increased dramatically for higher spatial frequencies and only slightly for lower spatial frequencies. Because inducing transparent motion between gratings of identical

spatial frequency is contingent on perceived differences in their respective depth planes, it follows that spatial frequencies would mitigate the disparity needed to induce transparent motion as we move away from the fixation plane.

One Final Point : High Spatial Frequency Constraints in Motion Revisited

In this thesis, we explored the literature surrounding high spatial frequency constraints on lower spatial frequency motion signals in the processing of broadband stimuli (e.g., Cleary & Braddick, 1990; Chang & Julesz, 1983; Eagle, 1998).

Although results from these experiments have been exploited as a robust effect, it is important to point out that some authors have considerably reduced these constraints even in the presence of high spatial frequencies. As pointed out by Eagle (1998), most ~~of~~ studies have used random-dot patterns. Such stimuli are problematic in that they possess a statistically flat spectrum. Consequently, the Fourier energy contained in any two-dimensional frequency band is proportional to $sf_h^2 - sf_l^2$ (where sf_h and sf_l represent the high and low spatial frequency cutoffs). This means that the energy at higher spatial frequencies is substantially greater than that at lower spatial frequencies. It has been argued that such stimuli are unnatural. In fact, spectral analyses of natural images conducted by Field (1987) have shown that spectra are not flat across spatial frequencies but tend to follow a $1/f^2$ spectrum. As a result, natural images tend to have equal apparent contrast across spatial frequency. This has consequences on results observed by Cleary and Braddick as flat spectrum stimuli may

have disproportionately stimulated channels tuned to higher spatial frequencies. The suggestion is that using a pattern with a $1/f^2$ spectrum may equilibrate responses across spatial scales by reducing the overall contribution of higher spatial frequencies. Thus, larger d_{max} values could be obtained using a $1/f^2$ spectrum.

Eagle (1997; 1998) tested this hypothesis by using broadband patterns spanning 6 octaves (0.23 to 15.0 cycles/degree), which either had a flat spectrum or a $1/f^2$ spectrum. The general observation across the flat and $1/f^2$ condition is that d_{max} values are greater in the $1/f^2$ spectrum condition. This suggests that for these stimuli, access to lower spatial frequencies was less affected by the presence of higher spatial frequencies.

It is interesting that broadband stimuli with a $1/f^2$ spectrum would exhibit larger d_{max} values. As argued by Eagle (1997; 1998), the increase in the d_{max} values occurs because of an even energy distribution across spatial scale. This contrasts with stimuli with a flat spectrum whose energy is disproportionately at higher spatial frequencies. Accordingly, high spatial frequency constraints observed by Cleary and Braddick may have been exaggerated due to the relative strength of responses of spatial frequency channels tuned to higher spatial frequencies. This argument is all the more interesting given that we have also observed substantial reductions in high spatial frequency constraints. However, in Experiments II and III of this thesis, the amplitude of constituent spatial frequencies in adjacent compound gratings were

equal. Thus, from a spectral perspective, our stimuli are closer in profile to those used by Cleary and Braddick. Yet high spatial frequency constraints were eliminated by simply introducing larger disparities. This is important as it illustrates that a change in the direction of interactions across spatial scales can be induced by larger disparities despite the presence of a strong signal from higher spatial frequencies.

It is tempting to draw parallels between the explanations offered for the reduction of high spatial frequency constraints in d_{max} and those observed for stereopsis in Experiment III of the current thesis. In fact, there are important similarities between the two proposed accounts. Both explanations suggest that high spatial frequency constraints occur because of relatively stronger responses in channels tuned to higher spatial frequencies. Also, both explanations offer a release from constraints through a reduction in the response of high spatial frequency channels. There is however an important difference as to how responses across spatial frequency channels are “calibrated” so as to favour signals arising from lower spatial frequencies. In d_{max} experiments, this is achieved by implementing a spectral profile that is favourable to lower spatial frequencies. In stereopsis, this is achieved by introducing larger disparities which are mainly processed by low spatial frequency channels.

Although there is no evidence that release from constraint due to manipulating spectral profile and/or disparity exercise their effects independently, it is important to

note that changing the spectral profile of a stimulus and changing the disparity between a pair of stimuli are very different manipulations. Indeed, manipulating the spectral profile of a stimulus mainly serves to establish the relative “visibility” of spatial scales in the stimulus, whereas disparity (as suggested in this thesis) serves mainly to specify which components should weigh more in correlating images across the two eyes. Thus, the possibility exists that they exercise their release from higher spatial frequency constraints independently, and that their effects are additive over a given range of performance. This would have important functional consequences as it would suggest that further increases in d_{max} could be observed beyond those obtained with a $1/f^2$ spectrum. This could potentially be assessed by comparing d_{max} values for flat and $1/f^2$ spectrum RDKs when presented binocularly at small and large disparities.

As a final note, it should be pointed out that the observation of a reduction in high spatial frequency constraints with the use of $1/f^2$ spectrum in d_{max} experiments could suggest that the constraints observed in experiment II simply resulted because our compound stimuli comprised a high and low spatial frequency of equal amplitude. This would have caused the perceived contrast at higher spatial frequencies to be higher than that at lower spatial frequencies. However, it must be noted that increases in d_{max} values for $1/f^2$ spectrum stimuli, do not represent a full release of high spatial frequency constraints. In fact, d_{max} values for $1/f^2$ spectrum stimuli still fall short (by approximately 25%) of the maximum d_{max} values observed

for low-pass stimuli. This indicates that high spatial frequency constraints are still present. Furthermore, Eagle (1998) demonstrated that constraints are greatest when high spatial frequency noise is adjacent to the signal spatial frequency and this even with a $1/f^2$ spectrum stimulus. This is a sensible result. In fact, although differences in contrast across spatial frequencies can be important over differences of two to three octaves for flat spectrum stimuli, these differences are considerably smaller across adjacent spatial frequencies.

Also, the task required in d_{max} experiments and the task required in Experiment II of this thesis are very different in nature. In d_{max} experiments observers are asked to report the direction of motion over increasingly larger displacements, a task which requires that observers access signals emanating from lower spatial frequencies. Given this, we should not be surprised that d_{max} values would increase with a stimulus in which high spatial frequencies are comparatively less visible ($1/f^2$ spectrum stimulus). In Experiment II of this thesis, observers were asked to resolve fine disparities, a task requiring that observers gain access to signals arising from higher spatial frequencies. From this, it is not clear whether or not using a $1/f^2$ spectrum would generate substantial reductions in the constraints observed in Experiment II. In fact, as argued in this thesis, if small disparities are mainly processed by high spatial frequency channels, the contribution of high spatial frequencies may still outweigh that of lower spatial frequencies in a $1/f^2$ spectrum

stimulus. Given this, although a reduction in the constraints observed in stereopsis can not be ruled out if a $1/f^2$ spectrum were to be implemented, a complete elimination of high spatial frequency constraints is unlikely to occur. Nonetheless, an experiment such as that proposed above where a flat and a $1/f^2$ spectrum stimuli are used at various disparities could resolve this issue and provide empirical data to further our understanding of the effects of disparity and spectral profile on spatial scale interactions in stereopsis at high velocities.

References

- Albrecht, D. G., and DeValois, R. L.. (1981). Striate cortex responses to periodic patterns with and without the fundamental harmonics. Journal of Physiology (London), 319, 497-514.
- Anderson, E. E. and Weymouth, F. M. (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.
- Anstis, S. M. and Harris, J. P. (1974). Movement aftereffects contingent on binocular disparity. Perception, 3, 153-168.
- Anzai, A., Ohzawa, I., and Freeman, R. D. (1999) Neural mechanisms for encoding binocular disparity: receptive field position vs. phase. Journal of Neurophysiology.
- Barlow, H. B., Blakemore, C., and Pettigrew, J. D. (1967). The neural mechanism of binocular depth discrimination. Journal of Physiology, 193, 338-354
- Blakemore, C. and Campbell, F. W. (1969). On the existence of neurons in the visual system selectively sensitive to the orientation and size of retinal images. Journal of Physiology (London), 225, 437-445.
- Bradley, D., Qian, N., and Andersen, R. (1995). Integration of motion and stereopsis in middle temporal cortical area of macaques. Nature (London), 373, 609-611.
- Campbell, F. W. and Robson, J. G. (1968). Applications of Fourier analysis to the visibility of gratings. Journal of Physiology (London), 197, 551-566.
- Chang, J. J., and Julesz, B. (1983). Displacement limits, direction anisotropy and direction versus form discrimination in random-dot cinematograms. Vision Research, 23, 639-646.
- Cleary, R., and Braddick, O. J., (1990). Masking of low frequency information in short-range apparent motion. Vision Research, 30, 317-327.

Cormack, L. K., Stevenson, S. B., and Landers, D. D. (1996). Correlation detection and utricular discrimination in displays with unequal monocular contrasts. Investigative Ophthalmology and Visual Science, supplement, 37, 652.

Cormack, L. K., Stevenson, S. B., and Landers, D. D. (1997). Interactions of spatial frequency and unequal monocular contrasts in stereopsis. Perception, 26, 1121-1136.

DeAngelis, G. C., Ohzawa, I., and Freeman, R. D. (1991). Depth is encoded in the visual cortex by a receptive field structure. Nature, 352, 156-159

DeAngelis, G. C., Ohzawa, I., and Freeman, R. D. (1995). Neural mechanisms underlying stereopsis: how do simple cells in the visual cortex encode binocular disparity? Perception, 24, 3-31.

Dev, P. (1975). Perception of depth surfaces in random-dot-stereograms: A neural model. International Journal of Man-Machine Studies, 7, 511-528.

Eagle, R. A. (1997). Independent processing across spatial frequency in moving broadband patterns. Perception, 26, 961-976.

Eagle, R. A. (1998). Upper displacement limits for spatially broadband patterns containing bandpass noise. Vision Research, 38, 1775-1787.

Felton, T. B., Richards, W., and Smith, R. A. (1972). Disparity processing of spatial frequencies in man. Journal of Physiology (London), 225, 349-362.

Field, D. J. (1987). Relationship between the statistics of natural images and the response properties of cortical cells. Journal of the Optical Society of America, 12, 2379-2391.

Foley, J. M. (1976). Binocular depth mixture. Vision Research, 16, 1263-1267.

Freeman, R. D. and Ohzawa, I. (1990). On the neurophysiological organization of binocular vision. Vision Research, 30, 1661-1676.

Glennerster, A. (1998). d_{max} for stereopsis and motion in displays. Vision Research, 38, 925-935.

Graham, N., Robson, J. G., and Nachmias, J. (1978). Grating summation in fovea and periphery. Vision Research, 18, 815-826.

Halpern, D. L. and Blake, R. (1998). How contrast affects stereoacuity. Perception, 17, 483-495.

Howard, H. J. (1919). A test for the judgment of distance. American Journal of Psychology, 2, 656-675.

Howard, I. P. and Rogers, B. J. (1995). Binocular vision and stereopsis. Oxford: Oxford University Press.

Hubel, D. H. and Livingston, M. S. (1987). Segregation of form, color and stereopsis in primate area 18. Journal of Neuroscience, 7, 3378-3415.

Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. Journal of Physiology, 160, 160-154.

Hubel, D. H. and Wiesel, T. N. (1970). Stereoscopic vision in macaque monkey. Nature, 225, 41-42.

Issa, N. P., Trepel, C., and Striker, M. P. (2000). Spatial frequency maps in cat visual cortex. Journal of Neuroscience (in press).

Julesz, B. (1960). Binocular depth perception of computer generated patterns, Bell System Technical Journal, 39, 1125-1162.

Julesz, B. (1971). Foundations of cyclopean perception. University of Chicago Press, Chicago.

Julesz, B. and Miller, J. E. (1975). Independent spatial frequency tuned channels in binocular fusion and rivalry. Perception, 4, 125-143.

Kaufman, L., Bacon, J., and Barroso, F. (1973). Stereopsis without image segregation. Vision Research, 13, 137-147.

Kwas, M., von Grunau, M., and Dubé, S., (1995). The effects of motion adaptation and disparity on motion integration. Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 49, 80-92.

Lankheet, M. J. M. and Lennie, P. (1995). Spatio-temporal requirements for binocular correlation in stereopsis. Vision Research, 36, 527-538.

Lappe, M. (1996). Functional consequences of an integration of motion and stereopsis in area MT of monkey extrastriate visual cortex. Neural Computation, 8, 1449-1461.

Legge, G. E., and Gu Y. (1989). Stereopsis and Contrast. Vision Research, 29, 989-1004.

Marr, D. and Poggio, T. (1976). Cooperative computation of stereo disparity. Science, 194, 283-287.

Marr, D. and Poggio, T. (1979). A computational theory of human stereo vision. Proceedings of the Royal Society of London, 204B, 301-328.

Maske, R., Yamane, S., Bishop, P. O. (1984). Binocular simple cells for local stereopsis: comparison of receptive field organisation for the two eyes. Vision Research, 24, 1921-1929.

Maunsell, J. H. R., and Van Essen, D. C. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. Journal of Neurophysiology, 49, 1127-1147.

Mayhew, J. E. W. and Frisby, J. P. (1981). Psychophysical and computational studies towards a theory of human stereopsis, Artificial Intelligence, 17, 349-385.

Morgan, M. J., and Castet, E. (1995). Stereoscopic depth perception at high velocities. Nature, 378, 380-383.

Odom, J. V. and Chao, G. M. (1987). A stereo illusion induced by binocularly presented gratings: Effects of number of eyes stimulated, spatial frequency, orientation, field size, and viewing distance. Perception and Psychophysics, 42, 140-149.

Parker, A. J. and Yang, Y. (1989). Spatial properties of disparity pooling in human stereo vision. Vision Research, 29, 1525-1538.

Pettigrew, J. D. Binocular interaction on single units of the striate cortex of the cat. Thesis (B. Sc.), University of Sydney, Australia, 1965.

Piggins, D. (1978). Moirés maintained internally by binocular vision. Perception, 7, 679-681.

Poggio, G. F. (1984). Physiological basis of stereoscopic vision. In Vision and vision dysfunction, Vol. 9, Binocular vision, (ed. D. Regan), pp.224-238. Macmillan, London.

Poggio, G. F. (1984). The analysis of stereopsis. Annual Review of Neuroscience, 7, 379-412.

Prince, S. J. D., Eagle, R. A., Rogers, B. J. (1998). Contrast masking reveals spatial-frequency channels in stereopsis. Perception, 27, 1345-1355.

Qian, N, Andersen, R. A., and Adelson, E. H., (1994). Transparent motion perception as detection of unbalanced motion signals. I. Psychophysics. Journal of Neuroscience, 14, 7357-7366.

Regan, D. and Beverley, K. I. (1973). Some dynamic features of depth perception. Vision Research, 33, 2359-2360.

Richards, W. J. (1951). The effects of alternating views of the test object on vernier and stereoscopic acuities. Journal of Experimental Psychology, 42, 376-383.

Robson, J. G. (1966). Spatial and temporal contrast-sensitivity functions of the visual system. Journal of the Optical Society of America, 56, 1141-1142,

Rohaly, A. M. and Wilson, H. R. (1993). Disparity averaging across spatial scales. Vision Research, 34, 1315-1325.

Schor, C., and Heckman, T. (1989). Interocular differences in contrast and spatial frequency : effects on stereopsis and fusion. Vision Research, 29, 837-847.

Schor, C., Heckman, T., and Tyler, C. W. (1989). Binocular fusion limits are independent of contrast, luminance gradient, and component phases. Vision Research, 7, 821-835.

Schor, C., Wood, I., and Ogawa, J. (1984). Binocular sensory fusion is limited by spatial resolution. Vision Research, 24, 661-665.

Shoham, D., Hubener, M., Schulze, S., Grinvald, A., and Bonhoeffer, T. (1997). Spatio-temporal frequency domains and their relation to cytochrome oxidase staining in cat visual cortex. Nature, 385, 529-533.

Simmons, K. (1984). Effects on stereopsis of monocular versus binocular degradation of image contrast. Investigative Ophthalmology and Visual Science, 25, 987-989.

Smallman, H. S. (1995). Fine-to-coarse scale disambiguation in stereopsis., Vision Research, 35, 1047-1060.

Smallman, H. S., and Macleod, D. I. A. (1994). Size-disparity correlation in stereopsis at contrast threshold. Journal of the Optical Society of America, 11, 2169-2183.

Smallman, H. S., and MacLeod, D. I. A. (1997). Spatial scale interactions in stereo sensitivity and the neural representation of binocular disparity. Perception, 26, 977-997.

Snowden, R.J., Treue, S., Erickson, R.E., and Andersen, R.A. (1991). The response of MT and V1 neurons to transparent motion. Journal of Neuroscience, 11, 2768-2785.

Taylor, M. M. and Creelman, C. D. (1967). PEST: Efficient estimates on probability functions. Journal of the Acoustical Society of America, 41, 782-787.

Tyler, C. W. (1971). Stereoscopic depth movement: Two eyes less sensitive than one. Science, 174, 958-961.

Tyler, C. W. (1973). Stereoscopic vision: Cortical limitations and a disparity scaling effect. Science, 181, 276-278.

Tyler, C. W., Schor, C. M., and Colletta, N. M. (1992) Spatio-temporal limitations of vernier and stereoscopic alignment acuity. Proceedings of the International Society for Optical Engineering. 1669, 112-121.

Van den Berg, A. V. and Brenner, E. (1994). Humans combine the optic flow with static depth cues for robust perception of heading. Vision Research, 34, 2153-2167.

von Grunau, M. W., and Dubé, S. (1993). Ambiguous plaids: Switching between coherence and transparency. Spatial Vision, 7, 199-211.

Waugh, S. J. and Hess, R. F. (1994). Suprathreshold temporal-frequency discrimination in the fovea and the periphery. Journal of the Optical Society of America A, 11, 1199-1212.

Wilson, H., Blake, R., Halpern, L., (1991). Coarse spatial scales constrain the range of binocular fusion on fine scales. Journal of the Optical Society of America A, 8, 229-236.

Woodburne, L. S. (1934). The effect of a constant visual angle upon the binocular discrimination of depth differences. American Journal of Psychology. 46, 273-286.

Appendix I

Table 1

Individual fitted χ parameters and standard errors as a velocity (deg/sec) and temporal frequency (Hz) (Exp. 1).

Observer		Spatial frequency (cycles/degree)					
		0.08	0.16	0.32	0.64	1.28	2.56
<i>Velocity (deg/sec)</i>							
SL	χ	412.5	189.2	81.0	29.3	14.0	4.4
	Std. Error	8.4	3.2	2.7	2.4	0.8	0.5
NC	χ	290.3	142.5	72.1	25.5	5.2	0.2
	Std. Error	4.1	1.0	1.6	2.4	1.0	0.8
LR	χ		137.5	70.9	24.6	10.2	2.5
	Std. Error		2.7	3.4	4.3	0.9	1.6
JG	χ		143.5	72.9	26.9	6.6	1.0
	Std. Error		5.6	1.9	3.8	2.5	0.4
<i>Temporal Frequency (Hz)</i>							
SL	χ	33.0	30.3	25.9	18.7	17.9	11.2
	Std. Error	0.7	0.5	0.9	1.5	1.0	1.2
NC	χ	23.2	22.8	23.1	16.3	6.7	0.5
	Std. Error	0.3	0.2	0.5	1.5	1.3	2.0
LR	χ		22.0	22.7	15.7	13.1	6.5
	Std. Error		0.4	1.1	2.8	1.2	4.1
JG	χ		23.0	23.3	17.2	8.5	2.6
	Std. Error		0.9	0.6	2.4	3.2	1.0