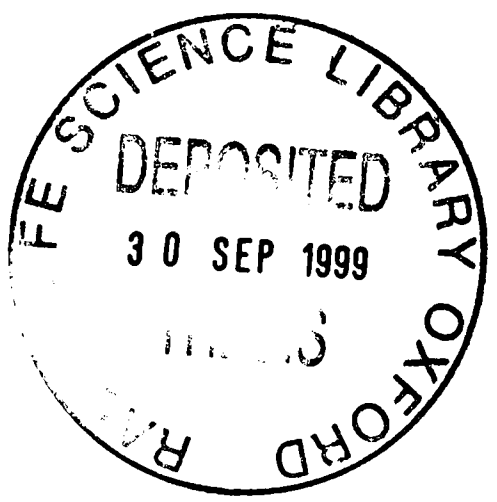
The Perception of Depth, Rotation and Shearing in Motion Parallax Surfaces

Sarah A. Swash



Thesis submitted for the degree of Doctor of Philosophy,

University of Oxford,

Trinity Term 1998.

Acknowledgements

I would like to express my warm thanks to Professor Brian Rogers, for his advice and help over the past few years and especially for his assistance in the production of the final version of this thesis. I also thank Mark Bradshaw, Richard Eagle, Maarten Hogervorst, Rachel Christophers, Alun Johns, Simon Prince and Matthew Lipson for their help and suggestions during my time in Oxford. Special thanks to Alun, Simon and Matthew for serving as subjects in this thesis.

Finally I thank my family, for their ongoing support, encouragement and sound judgment.

Sarah A. Swash
Lady Margaret Hall
October 1998.

To Mum, Dad, Edward and Emily

ABSTRACT

The Perception of Depth, Rotation and Shearing in Motion Parallax Surfaces

Thesis submitted by Sarah A. Swash of Lady Margaret Hall, Oxford,
for the degree of Doctor of Philosophy
during Trinity Term, 1998

Motion parallax is often considered to be an inherently ambiguous cue to depth. Despite the theoretical ambiguity associated with the pattern of retinal image motion, motion parallax generally evokes compelling three-dimensional (3-D) percepts and for this reason is regarded as an important source of 3-D information. Certain studies have indicated, however, that a parallax surface that contains a given amount of simulated depth is often perceived to rotate, rather than simply remain stationary, as the observer moves. This thesis provides an experimental investigation of the factors which influence perceived rotation and shearing in motion parallax surfaces.

In a series of psychophysical experiments, the cue of self-produced motion parallax was manipulated in order to provide insights into the mechanisms underlying the perception of 3-D surfaces. Since larger parallax motions often produce the impression of rotation, the "transition point" between stationarity and rotation was measured as a function of several factors. The maximum motion gradient was shown to be the principal determinant of this transition point — surfaces with a steep motion gradient were perceived to rotate at lower relative motion amplitudes than surfaces with shallow motion gradients. Vertical perspective information played a smaller role. The transition point also fell with increasing viewing distance. At even higher amplitudes, parallax surfaces can appear non-rigid or even lacking in all 3-D structure, and the experiments reported have measured the transition points between each of the different perceptual zones.

A model was introduced in order to determine whether the perceived magnitudes of depth and rotation of sinusoidal parallax surfaces were in accordance with geometric constraints. Qualitative support was found for a trade-off between depth and rotation when corrugation frequency or stimulus size was manipulated. Results were less conclusive when a dynamic vertical perspective cue was varied. A similar model was applied to the perceived magnitudes of depth and shearing in square wave surfaces. The relationship between these attributes was less clear-cut. The perceived rotation of sinusoidal surfaces increased with increasing viewing distance; possibly due to a decreasing propensity of the observer, with increasing viewing distance, to attribute the vertical perspective changes to self-motion.

In sum, these experiments demonstrate the importance of motion gradients and vertical perspective information in the perception of motion parallax surfaces, and suggest that the surfaces are generally perceived in qualitative accord with the totality of visual information present (Helmholtz, 1909).

Extended Abstract

The Perception of Depth, Rotation and Shearing in Motion Parallax Surfaces

Thesis submitted for the degree of Doctor of Philosophy
by
S.A. Swash, Lady Margaret Hall, University of Oxford,
during Trinity Term, 1998.

Psychophysical techniques were used to examine the percepts arising from motion parallax surfaces containing sinusoidal depth corrugations. Percepts of depth and rotation arising from head-linked parallax surfaces, and perceptions of depth order arising from surfaces in which the parallax was externally produced while the observer remained stationary, were experimentally investigated. A separate set of experiments examined the percepts of depth and shearing in square wave motion parallax surfaces.

A number of studies has indicated that motion parallax surfaces are perceived to rotate rather than simply translate as the observer or surface moves (Braunstein and Andersen, 1981; Ono, Rivest and Ono, 1986; Rogers and Collett, 1989; Ono and Steinbach, 1990). In this thesis, the amplitude at which a perceptually rigid, stationary parallax surface becomes perceptually rigid but rotating was investigated as a function of a variety of factors. Over a range of corrugation frequencies this 'Transition Point' was found to be largely determined by the maximum motion gradient (MMG) present in the surface. At higher frequencies the Transition Point became dependent upon the absolute relative motion amplitude. These results were confirmed and extended using 'clipped' corrugations in which the MMG was varied while the fundamental corrugation frequency remained constant.

For surfaces containing discontinuous motion gradients (such as ramp and square waveform gradients), the Transition Point function was best described by an analysis which takes into account the range of motion gradients in the surface. Increases in stimulus size resulted in increases in the Transition Point — larger surfaces were less likely to be seen as rotating — such that for each doubling of surface size the Transition Point increased by a factor of approximately 0.2. The Transition Points of proximally identical parallax stimuli viewed from 57, 114 and 228 cm decreased with increasing viewing distance — more distant surfaces were more likely to be seen as rotating — demonstrating that the rotation threshold is not quantifiable merely in terms of the proximal pattern of image motion. The slopes of the functions ranged between -0.5 and -0.7.

It has been suggested by a number of experimenters that the perception of different parallax surfaces can fall into different 'zones', which are separated by perceptual thresholds. Previous studies have indicated that there exist at least five such zones, although no single experiment to date has mapped them all. In this thesis four perceptual thresholds were determined as a function of corrugation frequency, in order to map five different perceptual zones. For a surface of given corrugation frequency, the zone into which the observer's percept falls can be manipulated through manipulation of the relative motion amplitude. The range of relative motion amplitudes resulting in a surface with given perceptual characteristics was found to vary with corrugation frequency. Surfaces of low corrugation frequency were perceived to be rigid and stationary over a relatively large range of amplitudes (8 - 9 arc min equivalent disparity). This range decreased as the surface corrugation frequency increased, such that surfaces of 1.6 cpd were perceived to be rotating as soon as the depth corrugations were visible. Conversely, the range of amplitudes resulting in a rigid, rotating percept was generally smallest for low frequency surfaces and largest for the mid-range of frequencies. The range of amplitudes supporting a non-rigid, rotating percept was high: depth was perceived in parallax surfaces containing amplitudes of 80 arc min equivalent disparity. This emphasises the propensity of the visual system to perceive relative motion as arising from a three-dimensional structure rather than from a two-dimensional source, even when the percept of depth is seen to be non-rigid.

The perception of rotation in the perceptual zone above the Transition Point was examined using a matching technique. The extent of rotation of a real, circular paddle situated above the display was adjusted so that it matched the perceived rotation magnitude of the test stimulus. It was found that the MMG played a smaller role in the determination of perceived rotation at these higher relative motion amplitudes than it had played in the determination of the Transition Point. Increasing the MMG through clipping the stimuli led to increased amounts of perceived rotation only for the surface of lowest corrugation frequency, 0.05 cpd. The perceived rotation of surfaces of 0.1 and 0.2 cpd remained unaffected by variations in the MMG. It was suggested that the perception of non-rigidity, at the expense of rotation, in the stimuli of higher corrugation frequency may partially account for this pattern of results.

Decreases in stimulus size resulted in increases in the magnitude of perceived rotation. Stimulus size played a larger role in perceived supra-threshold rotation than it played in the determination of the Transition Point. The reversal of the relative weightings of the motion gradient and stimulus size at supra-threshold rotation levels relative to threshold levels was suggested to be partly due to the comparative importance of the vertical perspective information (which signalled zero rotation) at supra-threshold levels of rotation.

Stimuli containing a central peak were perceived to rotate to a greater extent than those containing a central trough when the surface corrugation frequency was 0.05 cpd. This trend was reversed, and was slightly smaller, for 0.1 cpd surfaces. The sign of the central corrugation made no difference to the magnitude of perceived rotation of 0.2 cpd surfaces.

The perceived magnitude of rotation of stimuli containing two different corrugation frequencies, one in the central region of the stimulus and the other in the peripheral flanking regions, generally increased with increasing mean corrugation frequency. The central frequency was found to be more important than the flanking frequency in the determination of perceived rotation.

Later experiments of this thesis examined the perceived magnitudes of

both depth and rotation, and compared them to the values predicted by a geometric model. The model is based on a theoretical trade-off between depth and rotation such that, for a given amplitude of relative motion, an increase in the magnitude of one percept results in the concomitant decrease in the magnitude of the other percept (Rogers and Collett, 1989). Previous work using stereo-parallax surfaces is consistent with the notion that the visual system applies a simple geometric constraint when faced with discrepant motion parallax and disparity cues (Rogers and Collett, 1989; Williams, 1993). The observer generally perceives a surface whose depth is consistent with the disparity information; the motion parallax information is incorporated into the observer's percept as a rotation of the surface. The work of this thesis examined observers' perceptions of depth and rotation when motion parallax was the only cue.

A depth-rotation trade-off was observed empirically when the corrugation frequency of the surface was manipulated: increases in frequency led to increases in the extent of perceived rotation and to simultaneous decreases in the extent of perceived depth. The trade-off relationship was also observed when stimulus size was manipulated: increases in stimulus size resulted in decreased amounts of perceived rotation and increased amounts of perceived depth. This is consistent with the notion that, under these circumstances at least, the visual system obeys this simple geometric rule.

Despite the existence of this trade-off relationship, the combined magnitudes of the two percepts exceeded those predicted by the geometric model. Thus for a given magnitude of perceived depth of a given surface, the extent of perceived rotation should have theoretically been lower than observed, and vice versa. Three factors which possibly underlie the lack of quantitative accordance between observers' percepts and the predictions of the model were evaluated. The suggested factors were firstly, an over-estimation of the extent of 2-D relative dot motion (Sato, 1989; Ono and Steinbach, 1990); secondly, the possibility of a given amount of relative motion being perceived both as depth *and* as rotation, and thirdly, biases due to the technique for measuring either the depth or the rotation.

The magnitudes of perceived depth and perceived shearing of square wave surfaces were also examined and compared to those predicted by a

second geometric model. It was again found that the combined perceived extents of these magnitudes were greater than those predicted by the model, with observers frequently perceiving shearing magnitudes which were greater than the extent of 2-D relative dot motion. Evidence in favour of a trade-off of shearing and depth was less strong, and observers' propensity to interpret some displays (particularly those containing the higher corrugation frequency and higher relative motion amplitudes) as 2-D rather than 3-D suggests that square wave motion gradients do not support the perception of depth as readily as sinusoidal gradients. Furthermore, the qualitatively different percepts (depth and rotation) which arise from surfaces containing the latter gradient are more easily separable than the percepts of depth and shearing which arise from the square wave surfaces. This highlights once more the importance of the motion gradient in the determination of the percept evoked by motion parallax.

The depth-rotation trade-off relationship in sinusoidal surfaces was also investigated as a function of viewing distance. Depth and rotation judgements were made for proximally identical stimuli at viewing distances of 57, 114 and 228 cm. It was found that an increase in viewing distance resulted in increases in the magnitudes of both perceived rotation and perceived depth. At the nearest viewing distance, 57 cm, the motion parallax was interpreted largely as depth, the magnitude of which was consistently over-estimated. Relatively small magnitudes of rotation were perceived. At the furthest viewing distance, 228 cm, the emphasis had shifted such that the majority of the relative motion was interpreted as rotation. Depth magnitudes were under-estimated. It was suggested that the lack of depth constancy typically perceived for parallax surfaces at viewing distances beyond approximately 1 m (Ono, Rivest and Ono, 1986; Rivest, Ono and Saida, 1989) is partly due to the increased extents of rotation perceived at greater distances. Again, the combined magnitudes of the two characteristics exceeded those predicted by the model.

In order to determine whether the changes in perceived depth which occur concomitantly with variations in either corrugation frequency or size are due to the specific nature of these manipulations, dynamic vertical perspective was introduced as a cue to rotation. Despite the fact that the magnitude of perceived rotation was significantly affected by the vertical

perspective cue, the extent of perceived depth in the majority of surfaces was not greatly affected. This suggests that there may be a degree of independence of the two percepts, and that it is not an increase in perceived rotation *per se* which results in a decrease in perceived depth.

In further experiments observers carried out a shape judgement task in which either the relative motion amplitude or the corrugation frequency of a triangle wave surface was adjusted until the sides of the corrugations appeared to be at right-angles to each other. A dynamic perspective cue simulated rotation angles of 0°, +6.5°, +13°, -2° and -5°, where positively-signed angles indicate that the simulated direction of rotation was opposite to the direction of head movement and negatively-signed angles indicate that the simulated rotation direction was in the same direction as the head motion. For surfaces in which the simulated direction of rotation was in the opposite direction to the head motion, it was found that greater relative motion amplitudes and greater corrugation frequencies were required to make the surface appear right-angled as the magnitude of rotation simulated by the dynamic vertical perspective cue increased. These result are consistent with a depth-rotation trade-off. In keeping with previous experiments in this area, the actual relative motion amplitudes and corrugation frequencies at which the surface appeared right-angled were less than predicted, indicating that depth is over-estimated. The fact that rotation was also perceived in these 'right-angled' surfaces concords with the previous experiments of this thesis which demonstrated that the combined extent of the depth and rotation percepts is sometimes greater than predicted.

Surfaces for which the specified rotation direction was the same as that of the observer's head motion were not perceived to rotate in the expected direction. Rather, they appeared to rotate in the opposite direction and were slightly non-rigid. Clear patterns in the relative motion amplitudes or corrugation frequencies required for the perception of a right-angled surface were not found. Despite the geometric validity of rotation in a direction either the same as *or* opposite to the head movement, the visual system is biased towards perceiving the latter. Consideration of the depth magnitudes which would be predicted to accompany each rotation direction and the vertical perspective changes which occur at the eye suggest that again, the observer perceives the surface whose depth and

rotational characteristics are in least conflict with the totality of visual information present.

In order to determine the effect of vertical perspective upon the magnitude of perceived rotation as a function of viewing distance, observers made rotation matches for stimuli at distances of 57, 114 and 228 cm. The magnitude of rotation simulated by the vertical perspective cue was either 0° or 15°. It was found that the effect of the vertical perspective cue on perceived rotation magnitude remained approximately constant over distance, but that the absolute perceived magnitude increased. This pattern of results is consistent with the notion that the observer is decreasingly likely, with increasing viewing distance, to attribute vertical perspective changes as arising from self-motion. Vertical perspective changes which remain constant as viewing distance increases are increasingly likely to be interpreted as object rotation.

A final experiment investigated the perceived depth order of surfaces in which the motion parallax transformation was not accompanied by observer motion. The observer remained stationary. The depth order specified by the motion parallax information in this case is inherently ambiguous (c.f. KDE). The extent to which the perceived depth order could be influenced by a dynamic vertical perspective cue was examined. The perceived depth order of the central corrugation was continuously monitored throughout each 30 second trial. It was found that the proportion of time during which the perceived depth order was consistent with the vertical perspective information was dependent upon both the magnitude of simulated rotation and stimulus size.

This thesis aimed to provide a detailed analysis of the factors influencing the percepts evoked by motion parallax information. Despite the inherently ambiguous nature of this information, observers generally experience compelling three-dimensional percepts which remain constant over time. The characteristics of these percepts were shown to be particularly dependent upon the relative motion gradient and vertical perspective information. At low relative motion amplitudes, the former was shown to be particularly important in determining the threshold of perceived rotation. Vertical perspective information was the more important determinant of supra-threshold levels of perceived rotation.

The observed trade-off between the magnitudes of perceived depth and rotation is consistent with the notion that the visual system is governed, qualitatively at least, by simple geometric constraints.

Contents

Chapter 1: Introduction 1

1.1 Cues to depth 1
1.2 Transformations of the retinal image as cues to depth 3
1.2.1 Rotations about the X or Y axis 3
1.2.2 Rotations about the Z axis 4
1.2.3 Translations along the X or Y axis 4
1.2.4 Translations along the Z axis 4
1.3 Motion Parallax 6
1.3.1 Definition of the term, and its history 6
1.3.2 The formal similarity of motion parallax and binocular stereopsis 8
1.3.3 Absolute motion parallax 8
1.3.4 Relative motion parallax 10
1.3.5 The history of relative motion parallax as a depth cue 13
1.3.6 Absolute depth judgements and depth constancy 18
1.4 The similarity of motion parallax and stereo 22
1.4.1 Depth/curvature thresholds 22
1.4.2 Subthreshold summation 24
1.4.3 Between-cue threshold elevation 24
1.4.4 Induced effect 25
1.4.5 Contrast effects 26
1.5 The ambiguity of motion parallax 29
1.5.1 The Rigidity Assumption 29
1.5.2 Classification of non-rigid motions 30
1.5.3 The perception of non-rigid structures 30
1.5.4 The perception of non-rigidity in rigid structures 32
1.5.5 The perception of rigidity in non-rigid structures 35
1.5.6 The validity of the rigidity assumption 36
1.6 The ambiguity of rigid motions 36
1.6.1 The perceptual interpretation of motion parallax when combined with stereo 37
1.6.2 The perceptual interpretation of motion parallax when presented alone 42
1.7 Conclusion 45

Chapter 2: General Methods 46

2.1 Hardware 46
2.1.1 Motion Parallax Display 46
2.1.2 Calibration of Motion Parallax Display 47
2.1.3 Experimental Sequence Control 48
2.1.4 Vertical Image Size Changes 49
2.1.5 Calibration of the Vertical Image Size 50
2.1.6 Measurement of magnitude of perceived rotation, using a real paddle 53
2.1.7 Calibration of the paddle used for rotation matching 56
2.2 Software 58
2.3 Subjects 59

2.4 Psychophysical Procedures 59
2.4.1 Method of Constant Stimuli 59
2.4.2 Method of Adjustment 59
2.4.3 Note 60

Chapter 3: The Transition Point between a Rigid, Stationary surface and a Rigid, Rotating surface; Perceptual Zones 61

3.1 The concept of perceptual zones divided by thresholds 61
3.2 The 'motion threshold' 62
3.2.1 The effect of waveform type 62
3.2.2 The effect of relative motion amplitude 62
3.2.3 The effect of viewing distance 63
3.2.4 The effect of head movement velocity 63
3.2.5 The effect of type of head movement 64
3.2.6 The effect of observer movement 64
3.3 Experiments of this chapter 68
3.4 Methods 70
3.5 Experiment 1a: Transition Point as a function of Corrugation Frequency 71
3.5.1 Introduction 71
3.5.2 Stimuli 71
3.5.3 Results 71
3.5.4 Discussion 74
3.6 Experiment 1b: Transition Point as a function of Corrugation Frequency of Peak-clipped surfaces 77
3.6.1 Introduction 77
3.6.2 Stimuli 79
3.6.3 Results 80
3.6.4 Discussion 90
3.7 Experiment 2: Transition Point as a function of Waveform Type 92
3.7.1 Introduction 92
3.7.2 Stimuli 95
3.7.3 Procedure 95
3.7.4 Results 95
3.7.5 Discussion 99
3.8 Experiment 3: Transition Point as a function of Stimulus Size 101
3.8.1 Introduction 101
3.8.2 Stimuli 103
3.8.3 Results 103
3.8.4 Discussion 103
3.9 Experiment 4: Transition Point as a function of Viewing Distance 107
3.9.1 Introduction 107
3.9.2 Stimuli 107
3.9.3 Procedure 107
3.9.4 Results 107
3.9.5 Discussion 108
3.10 Experiment 5: Perceptual Zones into which the percept arising from a motion parallax stimulus can fall 110
3.10.1 Introduction 110
3.10.2 Stimuli 111
3.10.3 Procedure 112
3.10.4 Results 112
3.10.5 Discussion 113
3.11 Summary and Conclusions 115

Chapter 4: Perceived rotation of motion parallax surfaces in the perceptual zone above the Transition Point....118

4.1 Introduction....118
4.2 Experiment 6: The relative rotation of two stimuli of different corrugation frequencies....118
4.2.1 Stimuli....118
4.2.2 Procedure....119
4.2.3 Results....119
4.2.4 Discussion....122
4.3 Experiment 7a: Matched Rotation as a function of Corrugation Frequency....123
4.3.1 Apparatus....123
4.3.2 Stimuli....123
4.3.3 Procedure....124
4.3.4 Results....124
4.3.5 Discussion....130
4.4 Experiment 7b: Matched Rotation as a function of Corrugation Frequency of Peak-Clipped surfaces....133
4.4.1 Introduction....133
4.4.2 Stimuli....133
4.4.3 Procedure....134
4.4.4 Results....134
4.4.5 Discussion....146
4.5 Experiment 8: Matched Rotation as a function of Stimulus Form....149
4.5.1 Introduction....149
4.5.2 Stimuli....150
4.5.3 Procedure....150
4.5.4 Results....150
4.5.5 Discussion....151
4.6 Experiment 9: Matched Rotation as a function of Stimulus Size....156
4.6.1 Introduction....156
4.6.2 Stimuli....156
4.6.3 Procedure....156
4.6.4 Results....157
4.6.5 Discussion....160
4.7 Experiment 10: Matched Rotation as a function of Central and Flanking Corrugation Frequencies....161
4.7.1 Introduction....161
4.7.2 Stimuli....161
4.7.3 Procedure....162
4.7.4 Results....162
4.7.5 Discussion....172
4.8 Summary and Conclusions....172

Chapter 5: The trade-off between the perceived magnitudes of depth and rotation in observer-produced motion parallax surfaces....175

5.1 Introduction....175
5.1.1 Work involving non-quantitative judgements....175
5.1.2 Work involving quantitative judgements....177
5.2 Experiment 11: Matched Rotation and Matched Depth as a function of Corrugation Frequency....199
5.2.1 Stimuli....199

5.2.2 Procedure 199
5.2.3 Results 200
5.2.4 Discussion 204
5.3 Experiment 12: Matched Rotation and Matched Depth as a function of Stimulus Size 214
5.3.1 Stimuli 214
5.3.2 Procedure 215
5.3.3 Results 215
5.3.4 Discussion 219
5.4 General Discussion 223
5.5 Experiment 13: Matched Shearing and Matched Depth in Square Wave surfaces as a function of Corrugation Frequency 230
5.5.1 Introduction 230
5.5.2 Stimuli 230
5.5.3 Procedure 231
5.5.4 Results 231
5.5.5 Discussion 237
5.6 Experiment 14: Matched Shearing and Matched Depth in Square Wave surfaces as a function of Stimulus Size 237
5.6.1 Stimuli 237
5.6.2 Procedure 237
5.6.3 Results 237
5.6.4 Discussion 242
5.7 Experiment 15: Matched Rotation and Matched Depth in Proximally Identical Stimuli as a function of Viewing Distance 243
5.7.1 Introduction 243
5.7.2 Stimuli 244
5.7.3 Procedure 244
5.7.4 Results 245
5.7.5 Discussion 261
5.7 Summary and Conclusions 262

Chapter 6: The Effect of Dynamic Vertical Perspective on Perceived Rotation and Perceived Depth in Motion Parallax Surfaces 271

6.1 Introduction 271
6.2 Experiment 16: The Effect of Dynamic Vertical Perspective on Matched Rotation and Matched Depth as a function of Corrugation Frequency 271
6.2.1 Apparatus 271
6.2.2 Stimuli 271
6.2.3 Procedure 272
6.2.4 Results 273
6.2.5 Discussion 279
6.3 Experiment 17: The Effect of Dynamic Vertical Perspective on Matched Rotation and Matched Depth as a function of Stimulus Size 280
6.3.1 Stimuli 280
6.3.2 Procedure 280
6.3.3 Results 281
6.3.4 Discussion 286
6.4 General Discussion 287
6.5 Experiment 18: The Effect of Dynamic Vertical Perspective on Shape Judgements: Task One 292
6.5.1 Introduction 292

6.5.2 Stimuli....292
6.5.3 Procedure....292
6.5.4 Results....293
6.5.5 Discussion....299
6.6 Experiment 19: The Effect of Dynamic Vertical Perspective on Shape Judgements: Task Two....302
6.6.1 Introduction....302
6.6.2 Stimuli....303
6.6.3 Procedure....303
6.6.4 Results....303
6.6.5 Discussion....304
6.7 Experiment 20: The Effect of Dynamic Vertical Perspective on Matched Depth and Matched Rotation as a function of Viewing Distance....309
6.7.1 Introduction....309
6.7.2 Stimuli....310
6.7.3 Procedure....310
6.7.4 Results....311
6.7.5 Discussion....318
6.8 Experiment 21: The Effect of Magnitude of Dynamic Vertical Perspective and Stimulus Size on Perceived Depth Order in Externally-Produced Motion Parallax Surfaces....320
6.8.1 Introduction....320
6.8.2 Stimuli....328
6.8.3 Procedure....328
6.8.4 Results....329
6.8.5 Discussion....335
6.9 Summary and Conclusions....336

Chapter 7: Conclusions....339

7.1 Thesis Aims....339
7.2 Summary....340
7.2.1 Chapter 3....340
7.2.2 Chapter 4....345
7.2.3 Chapter 5....348
7.2.4 Chapter 6....352
7.3 Discussion....356

Appendix A....359
Appendix B....364
Appendix C....365
Appendix D....371

References....373

Chapter 1:

Introduction

1.1 Cues to depth

The perception of a three-dimensional world with eyes which have two-dimensional projection surfaces has in the past been regarded as a paradox. Porta, in 1589 (referenced in Braunstein, 1976), made an analogy between the eye and the camera obscura — an apparatus that uses a darkened room (or box) with an aperture for projecting the outside scene onto the inside surface of the wall opposite the hole. Early astronomers used the camera obscura to watch the solar eclipse. The analogies which were drawn by Porta between the physical workings of the camera and the eye were generally legitimate, but difficulties arise when the question of how the owner of the eye actually gains an accurate impression of his surroundings is posed. Since the projection surface of the eye — like the camera — is flat, early empiricist philosophers concluded that depth could not be directly perceived. Locke (1694) and Bishop George Berkeley (1709) concluded that it had to be inferred from visual sensations. Upon receipt of the two-dimensional images of the surrounding visual scene, the mind (or some undefined internal process) proceeded to reconstruct the three-dimensional structure of the viewed scene by using the various 'cues' to depth in the two-dimensional images and in the observer's sensations arising from the muscles controlling the eyes. Such extreme theories of indirect visual perception have now, to an extent, been succeeded. The term 'cue' has persisted, although it no longer seems to have such strong theoretical connotations concerning the processes which result in the perception of a three-dimensional world.

Accommodation and convergence have in the past been regarded as important cues to depth. When a distant object is fixated, the lines of sight are almost parallel and the angle between these, the convergence angle, is small. Conversely, when the object is near, the convergence angle is large.

Descartes (1937; referenced in Howard and Rogers, 1995), ".. described the eyes as 'feeling out' a distance by a convergence of their optic axes, just as a blind man might feel out a distance with two staves" (Howard and Rogers, 1995, p. 427). Convergence was one of Berkeley's (1709) three proposed means of indirect perception of depth, the other two being 'confusion' — the blurring of an object brought close to the eye — and 'straining', the effort required to focus on a nearby object. Both of these are consequences of accommodation, the adjustment of the lens curvature required to bring an object into focus.

Accommodation and vergence, along with binocular disparity (see Section 1.3.2), were once regarded as the primary sources of depth information. Graham (1965, p. 504) suggested that convergence is mainly used in coarse near-far discriminations within a few yards of the observer. Judgements of relative depth have been shown to be more accurate when observers shift fixation between a nearer and a further object than when they fixate one of the objects only (Foley and Richards, 1972). The further target in this study was at a viewing distance of 250 cm. Since convergence does not change sizeably for viewing distances beyond approximately 2 m, it seems likely that this source of depth information is restricted to fairly close distances.

Since a picture can also give rise to vivid depth percepts in the presence of accommodation, convergence and disparity cues which indicate a flat surface, a number of 'secondary' or 'pictorial' cues to depth have also been defined. Two examples of monocular, static depth cues are texture gradients and shading. The former cue works on the assumption that a given surface has a texture of uniform density, so that when the surface is projected to the eye the regions of the surface which are further from the observer will have a higher texture density on the retina than the nearer regions. Gibson and Bridgeman (1987) have shown that textural information provides cues to depth and distance in photographs. Shading refers to the location and nature of the pattern of illumination on the object's surface, arising from the fact that the object is illuminated from a given location. Areas of highlights and shading have been shown to provide depth information in two-dimensional pictures (Berbaum, Bever and Chung, 1983, 1984; Berbaum, Tharp and Mroczek, 1983).

1.2 Transformations of the retinal image as a cue to depth

An important set of depth cues are those which involve transformations of the retinal images over time. These transformations arise as a result of movement either of the observer or of parts of the viewed scene. An infinite variety of motions can occur in three-dimensional (3-D) space. Braunstein (1976) has defined a subset of six types of motion which have commonly been studied in perceptual research. These involve stimulus rotation about and translation along each of three orthogonal axes: the Z axis, which is the observer's line of sight, the X axis, which is parallel to the line joining the two eyes and perpendicular to the Z axis and the Y axis, which is vertical and orthogonal to the X and Z axes. Through the consideration of *observer* translations and rotations, Rogers (1993) defined a further six categories of motion.

1.2.1 Rotations about the X or Y axis

The retinal image transformations resulting from object rotations about either the X or the Y axis provide information about the object's 3-D structure. Miles (1929,31) demonstrated that the shadows of rotating objects give rise to the perception of rigid motion, but his work did not specifically explore the conditions which elicited this percept. Wallach and O' Connell (1953) were the first to experimentally investigate these transformations in the optic array using a shadow projection technique similar to that of Miles. Various objects (solids, wire figures, straight rods) were placed between a point light source and a translucent screen, onto which the shadows of the objects were projected. The distance between the light source and screen was large, making the projection effectively parallel. Static shadows did not appear 3-D. Rotation of the object evoked a compelling 3-D percept, thus allowing the isolation of the effects of motion on perceived depth. Wallach and O' Connell named this effect the Kinetic Depth Effect (KDE).

The transformations in the optic array which result from observer rotations about the X or the Y axis (e.g. as produced by vertical or horizontal eye or head movements) do not provide information about the 3-D structure of the environment. The optic flow potentially provides

expropriospecific information concerning the magnitude and direction of the eye or head movements (Rogers, 1993).

1.2.2 Rotations about the Z axis

The image transformations arising from object rotation about the Z axis do not provide depth information, because all parts of the object rotate about the axis with the same angular speed. As pointed out by Rogers (1993), the Stereo-Kinetic Effect (SKE) is *not* a retinal image transformation resulting from object rotation about the Z axis (Braunstein, 1976). Rather, SKE demonstrations are approximations of the transformations resulting from the partial rotation or 'wobbling' of a structure about its X and Y axes.

Transformations resulting from rotation of the observer about the Z axis (e.g. torsional eye or head movements) can not provide depth information, for the reason in the preceding paragraph. The patterns may, however, provide expropriospecific information about eye or head torsion (Rogers, 1993).

1.2.3 Translations along the X or Y axis

Motion parallax is the optic array transformation which accompanies translations along the X or the Y axis. Most germane to this thesis are the transformations along the former axis. Situations in which the observer translates are cases of 'observer-produced parallax', and those in which the object translates are cases of 'object-produced parallax'. Motion parallax is considered in greater detail in Section 1.3 below.

1.2.4 Translations along the Z axis

Optic array transformations of object translations along the Z axis were studied by Kilpatrick and Ittelson (1951). Their apparatus allowed both simulated and real translations. Simulated translations were effected through manipulation of the size of the stimulus in the frontal plane. The apparatus was a light box mounted behind a variable diaphragm. The diaphragm could be moved along tracks parallel to the observer's line of sight. The size of the diaphragm and its distance from the observer could be independently manipulated so that as it approached the observer its

projected size could be made to increase, remain constant, or decrease. Monocular observers perceived the stimulus to approach in the first case, remain stationary in the second and recede in the third.

It has been suggested that responses to contracting and expanding patterns of optic flow are innate. Eight-day old babies show defensive stress reactions when a foam rubber cube is pushed towards them (Bower, Broughton and Moore, 1970) or when viewing optical displays depicting expansion patterns (Bower et al., 1970; Ball and Tronick, 1971). The same reactions have been found in infant rhesus monkeys (Schiff, Caniness and Gibson, 1962). Regan and Beverley (1978) suggested the existence of a 'looming detector' which is specifically sensitive to contracting and expanding patterns of optic flow.

More recent work has demonstrated that the expanding optic flow pattern which accompanies an approaching frontal surface can theoretically be used to compute the estimated time to contact (Lee, 1980; Lee and Young, 1985). There is evidence suggesting that humans can indeed make accurate estimates of time to contact (e.g. Schiff and Detwiler, 1979; Freeman, Harris and Tyler, 1994).

As noted by Rogers (1993), the above categories do not classify the optic array transformations themselves, but rather the ways in those transformations can be produced. It is not necessarily the case that 3-D structure can be recovered from the transformations. Moreover, combinations of various motions within the above categories might give rise to identical optical patterns. Several mathematical analyses of optic flow have been produced. Optic flow patterns emphasise the importance of the totality of optical motions, relative and absolute, induced by observer motion (Gibson, 1950). Flow patterns provide information about both the motion of the observer and the structure of the environment (Gibson, 1950, 1966). Gordon (1965) provided one of the first formal analyses of the instantaneous velocity field. Koenderink and van Doorn (1975, 1976) and Longuet-Higgins and Prazdny (1980) also provided formal analyses and classified local changes in velocity into (i) expansion (dilatation), (ii) rotation (vorticity) and (iii) deformation (shear). Deformation, which arises during parallax transformations, can provide information about the local slant of a surface up to a scale factor of

viewing distance. Nakayama and Loomis (1974) described a simple physiological mechanism which could analyse global patterns of optic flow. They defined a higher-order variable of optic flow, 'convexity' which, under the assumption of a rigid environment, is shown to be related to relative depth. Through the linkage of velocity-sensitive cells in a centre-surround organisation, the authors defined a higher-order cell, the 'convexity' cell, which is sensitive to discontinuities in the flow field. Layers of such cells could therefore provide a structuring of the visual field into distinct surfaces. The cells are insensitive to rotations, and consequently respond in an invariant manner to the aspects of optic flow which arise due to the structure of the environment.

1.3 Motion Parallax

1.3.1 Definition of the term, and its history

The term 'parallax' comes from astronomy, and is the difference in direction of a body resulting from a difference in the position of the observer. The parallax of a star is obtained by measuring its apparent displacement relative to more distant stars when the earth is in two different positions: to the right of the sun (in mid-winter) and to the left of the sun (in mid-summer). From its parallax, the distance of the star can be calculated. A parallax value of one second of arc corresponds to a distance of one 'parsec', 200,000 times the distance to the sun (Dyson, 1929).

The term 'monocular motion parallax' stems from the original astronomical term, and it came into use when the study of visual space perception began (Gibson, Olum and Rosenblatt, 1955). Motion parallax arises as a result of the movement of either an observer or an object (or both). It is the relative movement of images across the retina (or other projection surface) produced by the 3-D structure of the viewed scene when there is relative motion between that scene and the observer.

A historical development of the use of motion parallax as a guide to distance is provided by Boring (1942). In this review it is pointed out that there seems to be no recognition of this particular use of motion parallax prior to the time of Helmholtz, possibly because it was considered self-

evident. Helmholtz (1925) described the sensation of near objects 'gliding past us as we advance through the countryside, of more distant objects doing so more slowly, and of remote bodies maintaining their permanent position in the field' (Gibson et al., 1955). A commonly observed instance of motion parallax is the apparent movement of the scenery when one is looking from the window of a moving train, where the above description seems particularly apt. Indeed, Gibson et al. (1955, p. 373) suggested that ".. relative apparent motion in the visual field probably became obvious when railroad trains made rapid locomotion a commonplace".

At this point it is useful to draw a distinction between two types of information given by motion parallax. Theoretically, parallax can specify both the *absolute* distance of a given point, and the *relative* distance between two points separated in depth. Helmholtz's descriptions of motion parallax suggested that both types of information could be gained. Firstly, he observed that as we move through the countryside,

"... the apparent angular velocities of objects in the field of view will be inversely proportional to their real distances away; and consequently safe conclusions can be drawn as to the real distance of the body".

Helmholtz (1925, p.295)

On the next page, he notes that a monocular observer standing in a dense wood would have difficulty in perceiving the true 3-D structure of the environment. However,

"... the moment he begins to move forward, everything disentangles itself, and immediately he gets an apperception of the material contents of the woods and their relations to each other in space, just as if he were looking at a good stereoscopic view of it."

(Helmholtz, 1925, p.296)

In the former quotation, Helmholtz is referring to absolute motion parallax, and is maintaining that the angular velocity of an object can be used to deduce its distance from the observer. In the latter quotation he is referring to relative motion parallax, and he is merely stating that parallax provides information concerning the relative depths of objects. Absolute motion parallax and relative motion parallax are discussed in Sections

1.3.3 and 1.3.4, respectively.

1.3.2 The formal similarity of motion parallax and binocular stereopsis

A stationary observer viewing a scene with two eyes will receive two slightly different views of that scene, due to the horizontal displacement of one eye relative to the other. The term 'binocular disparity' refers to the differences between these views. Taken in its most commonly-used sense, it refers to the differences in the horizontal positions of image features between the two eyes — relative horizontal disparity. This is the sense in which the term will be used in this thesis. Other types of disparity, for example vertical disparity (Longuet-Higgins, 1982; Mayhew and Longuet-Higgins, 1982; Gillam and Lawergren, 1983), also exist. For a detailed discussion of the various types of disparity, see Howard and Rogers (1995).

Motion parallax and binocular disparity are formally equivalent under certain conditions, and it has been suggested that the same laws may apply to both (e.g. Barnard and Thompson, 1980; Gogel, 1978; Rogers and Graham, 1982). Helmholtz (1866) was the first to note the general similarity between parallax and stereopsis. He wrote

> *"... since the two eyes occupy positions in space that are not quite the same, the objects in front of us are seen from two slightly different points of view.. there is the same kind of difference in the images as would be produced by moving in space from one place to the other."*
>
> (Helmholtz, 1925, p. 295)

When the observer views a stationary object whilst translating horizontally through the interocular distance, the geometry which governs the motion of the object's projected image is formally identical to that which governs binocular disparity. This is illustrated in Figs. 1.1 and 1.2. The most important difference between the two situations is that in the parallax case, the two views are separated in time, whereas in the stereo case the two views are obtained simultaneously.

1.3.3 Absolute motion parallax

Absolute motion parallax involves only a single point. The absolute

distance between the observer and the fixated point can theoretically be recovered through measurements of the amount of rotation of the eye during a lateral head movement, and the lateral distance travelled by the eye. The absolute distance under binocular viewing is recoverable by measuring the vergence angle and the interocular distance. Fig. 1.1 illustrates the analogy between these situations.

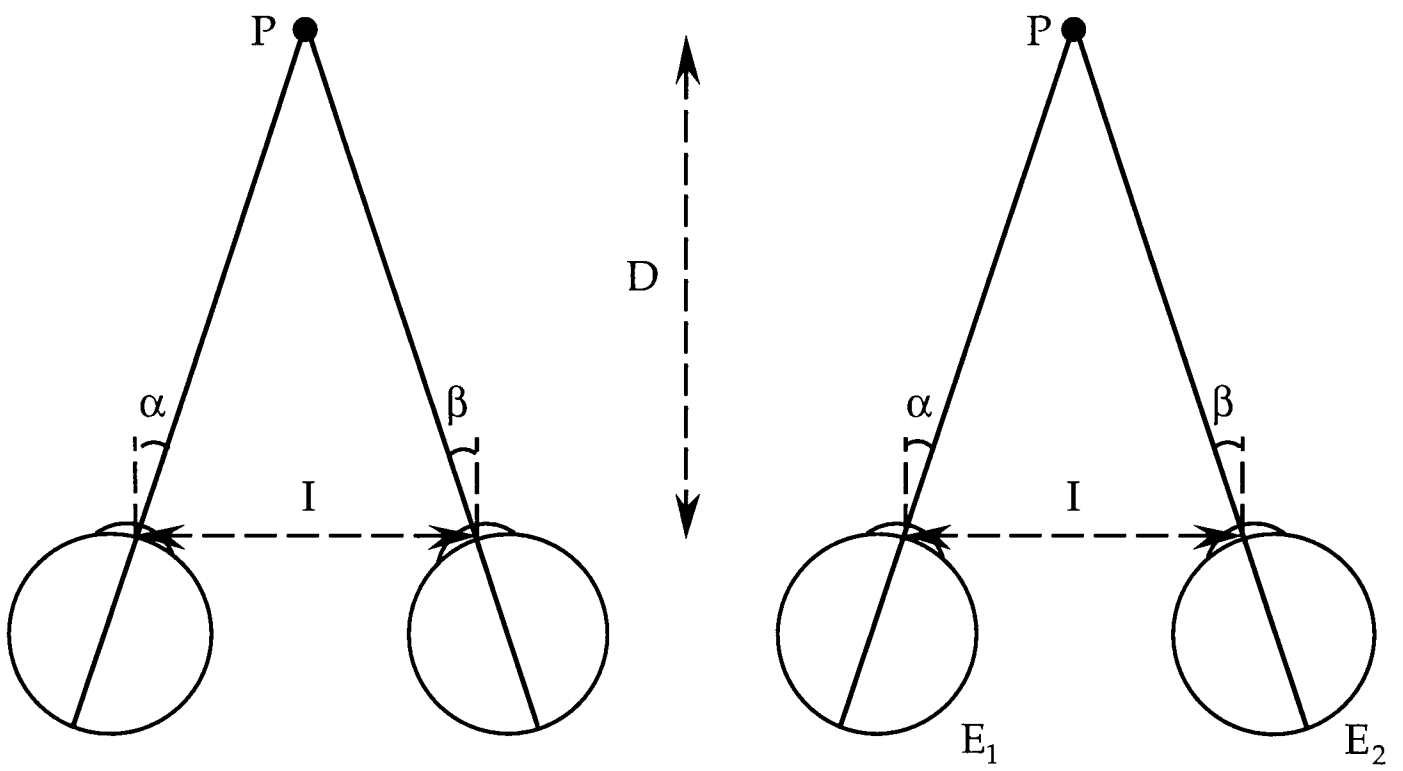


Fig. 1.1: A. Absolute Disparity. Knowledge about (i) the visual directions of point P in two simultaneous views and (ii) the separation between the eyes, is theoretically sufficient to obtain distance D.
B. Absolute Motion Parallax. The same information is provided when the eye translates through the interocular distance, from E_1 to E_2.

Although absolute parallax can theoretically provide the information necessary for the recovery of the absolute distance between object and observer, empirical studies have provided only limited support for its efficacy. Neither Park (1964) nor Bourdon (1902, cited in Gibson, Gibson, Smith and Flock, 1959, p. 41) found evidence for its use. Eriksson (1974, Experiment 1) presented observers with figures at distances of 200, 450 and 700 cm. The figures were constructed so that those at the furthest distance had the largest proximal size and those at the nearest distance had the smallest. Under stationary viewing conditions the figures appeared to lie in the same plane, but upon observer movement the relative depths of the figures were accurately perceived. Distances to each set of figures were, however, underestimated.

Gogel (1969) found that when observers viewed a light under conditions of reduced distance cues whilst translating their head horizontally, inaccurate estimations of the light's distance were made. The errors generally tended to place the light at a distance of approximately 2 m. This has been labelled the 'specific distance tendency'. Gogel and Tietz (1973) further investigated this bias, and asked subjects to estimate the distance of the light and its perceived extent of motion when head translation occurred. Obviously, if the light remains physically stationary and its distance from the observer is correctly perceived, it will (in theory) appear to remain stationary (although see Shebilske and Proffitt, 1981; 1983). It was found once again that distance estimations tended to 2 m and that the extent of the light's apparent motion varied with its actual distance from the observer. These results were based on four back-and-forth movements of the head. When this number was increased, the amount of apparent motion of the light decreased. Distance estimates were not taken in this latter instance however, and it can not be concluded that the lack of apparent motion would necessarily have led to an accurate distance estimate. On the basis of these findings the authors suggested that motion parallax was ineffective as a cue to absolute distance.

Both Dees (1966) and Ferris (1972) also provide evidence that motion parallax does not supply distance information for inexperienced observers but that performance can be markedly improved via training, a finding which fits with the study by Gogel and Tietz (1973). The evidence reviewed suggests that untrained subjects are generally unable to provide accurate distance estimates on the basis of absolute parallax.

1.3.4 Relative motion parallax

Relative motion parallax provides information concerning the depth between objects, or between two points on a single object. In principle, since absolute distance information is theoretically provided by the absolute parallax of each point, the absolute depth within the object is also available. The same information is theoretically provided by relative disparity. The meaning of these terms is illustrated below (Figs. 1.2 and 1.3).

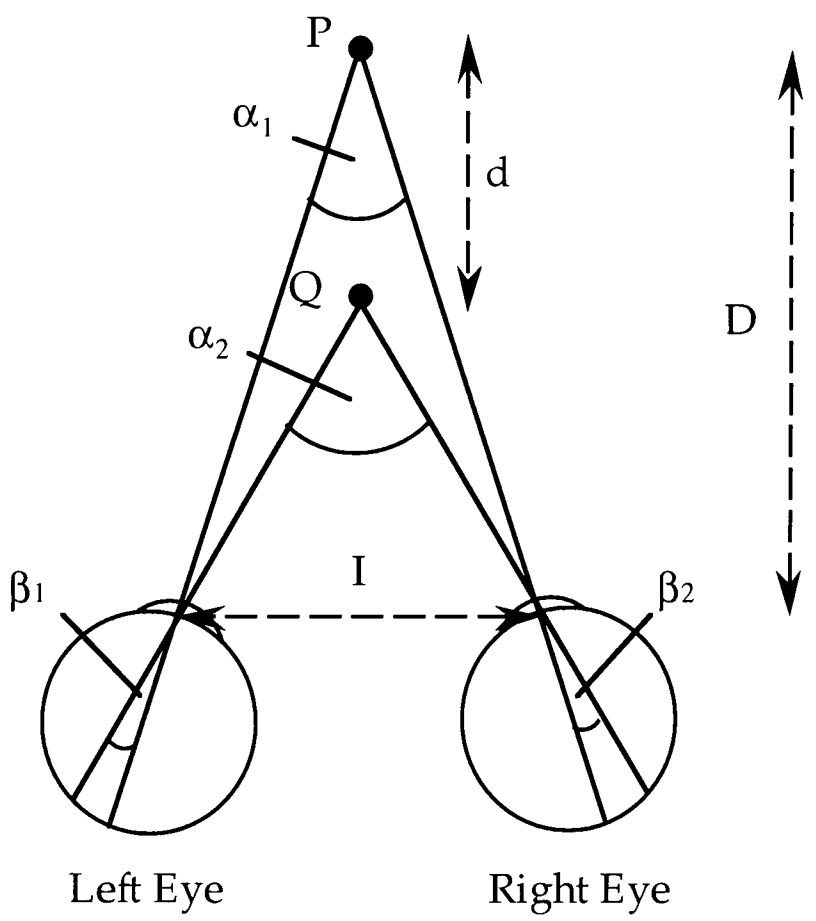


Fig 1.2: Geometry of the binocular viewing of a fixated point P and a non-fixated point Q.

In Fig. 1.2 above, the observer fixates Point P. Point Q lies closer to the observer by a distance of d. The light ray from Point Q falls to the right of the fovea in the right eye, and to the left of the fovea in the left eye. The disparity (in radians) between the images of the two points is given by the difference between β_1 and β_2 (which are oppositely signed), β_1-β_2.

$$\beta_1 - \beta_2 = \alpha_2 - \alpha_1.$$

Since D is large relative to I and d, $\tan \alpha_1 \approx \alpha_1$ and $\tan \alpha_2 \approx \alpha_2$, thus

$$\alpha_1 \approx I/D \text{ and } \alpha_2 \approx I/(D-d).$$

Thus the disparity, $\alpha_2 - \alpha_1 = Id/(D(D-d))$, which approximates to Id/D^2.

Fig. 1.3 illustrates the application of this geometry to motion parallax.

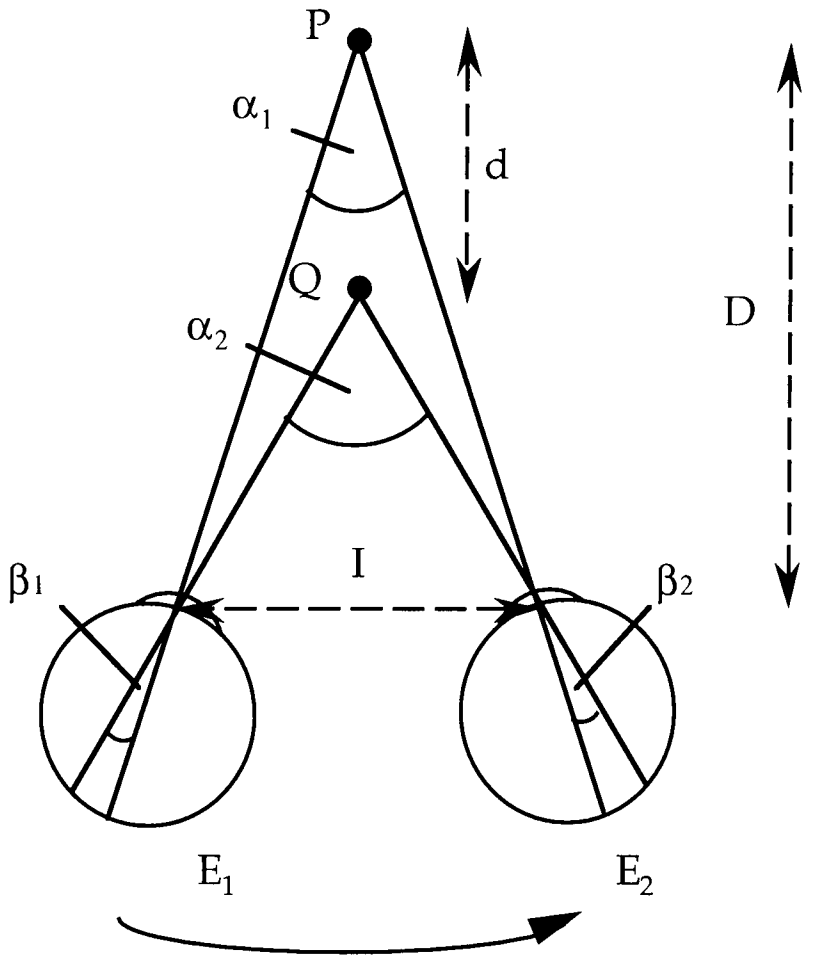


Fig 1.3: Geometry of the motion parallax between a fixated point P and a non-fixated point Q.

The observer fixates point P. Point Q lies closer to the observer by a distance of d. When the eye is at position E_1 the light ray from Point Q falls to the left of the fovea. As the eye moves from E_1 to E_2 there is relative motion between the images of P and Q such that when the eye reaches E_2 the ray from Q falls to the right of the fovea. The relative motion (in radians) between the images of the two points is given by the difference between β_1 and β_2 (which are oppositely signed), β_1-β_2.

$$\beta_1 - \beta_2 = \alpha_2 - \alpha_1.$$

Since D is large relative to I and d, $\tan \alpha_1 \approx \alpha_1$ and $\tan \alpha_2 \approx \alpha_2$, thus

$$\alpha_1 \approx I/D \text{ and } \alpha_2 \approx I/(D-d).$$

Thus the relative motion, $\alpha_2 - \alpha_1 = Id/(D(D-d))$, which approximates to Id/D^2.

The amount of relative motion for a lateral translation of the eye through the interocular distance (typically taken to be 6.5 cm), has been termed 'equivalent disparity' (Rogers and Graham, 1982).

Despite the theoretical availability of absolute depth from the motion parallax of two points, Section 1.3.3 suggests that observers can *not* extract absolute distance information from absolute parallax (Bourdon, 1902; Park, 1964; Dees, 1966; Gogel, 1969; Eriksson, 1974; Ferris, 1972; Gogel and Tietz, 1973). Consequently, the depth information in situations involving relative motion parallax is available only up to a scale factor. However, alternative sources may specify absolute distance. Evidence for depth constancy in motion parallax surfaces is discussed in Section 1.3.6.

As noted by Rogers (1993), the extraction of relative depth between points is simpler than the extraction of absolute depth (Koenderink and van Doorn, 1975), thus it might be expected that percepts of this nature might be more compelling than those related to the absolute distance of an object, and consequently that judgements involving relative depth will be more accurate than those involving absolute depth. Motion parallax was defined above as the relative movement of images across the retina (or other projection surface) produced by the 3-D structure of the viewed scene when there is relative motion between that scene and the observer. As noted above, the type of parallax most pertinent to this thesis involves observer/object translations along the X axis, and when the term 'parallax' is used in the experimental literature, it almost always refers to this category.

1.3.5 The history of relative motion parallax as a depth cue

Geometrical considerations show that motion parallax can theoretically provide information both about the absolute distance between the observer and a given object, and about the relative depths within the object itself. The notion that monocular depth cues can be perfectly adequate for the perception of 3-D space is supported by the medical observation that the gradual loss of the sight of one eye can go unnoticed (Metzger, 1966). One-eyed children can make effective use of motion parallax (Gonzalez, Steinbach, Ono and Wolf, 1989). Insurance adjusters have recognised the utility of parallax in the disambiguation of stimuli in different depth planes, and take this into consideration when dealing with claims following the loss of an eye (Hell, 1981).

These studies, along with the observation of Julesz (1971) that

approximately 15% of observers have difficulties in perceiving depth in complex random dot stereograms and that approximately 2% are completely stereo blind, but that comparable 'parallax blindness' is extremely rare, highlight the importance of motion parallax as a depth cue. 'Motion blindness' has been observed in a much studied individual, LM, who experienced selective disturbance of movement vision following a venous thrombosis which produced bilateral lesions of the temporo-occipital cortex (Zihl, Von Cramon and Mai, 1983). Motion perception in both horizontal and vertical dimensions was affected. The patient saw the world in terms of a series of 'snapshots', especially if movement was swift; she had trouble pouring a cup of tea, for example, because she was unable to see movement but rather perceived the liquid as 'frozen' and thus could not determine when to stop pouring.

It has long been recognised that motion parallax can aid visual space perception, although early experimental work did not prove that this was definitively the case. Such work (e.g. Bourdon, 1902; Cords, 1913; Tschermak-Seysenegg, 1939; Graham, Baker, Hecht and Lloyd, 1948; Zegers, 1948) seemed to indicate that the cue could indeed be used to make judgements about depth within the displays, but did not *isolate* motion parallax as a depth cue. In each of the procedures there were additional depth cues which may have, to varying extents, aided the observer in their performance. Bourdon (1902) showed observers two luminous spots at different distances in a dark corridor. The spots had the same proximal size. Viewing was monocular. When the observer's head was fixed in position with a biteboard, subjects were unable to judge which spot was nearer. However, upon 'the slightest movement' of their head from side to side, relative depth was easily judged. The absolute distance was not detectable.

Tschermak-Seysenegg (1939) used a 'parallactoscope' to examine the perceived relative depth between two objects. Subjects moved their head from side to side on a sliding headrest and equated the distance of two vertical wires. Performance was better in this condition than when subjects viewed the wires while remaining stationary, but was not as good as performance in a condition in which two eyes were used.

Rock (1984) carried out experiments in which observers viewed nine

small luminous discs, which were mounted on screens at three different viewing distances. Side-to-side head movements did not allow observers to make accurate judgements of the relative depths of the objects, leading to Rock's conclusion that "... motion parallax does not by itself seem to be a cue to distance or depth" (although see Rock, 1995).

In the study by Graham et al. (1948), it was reported that some subjects' performance was based merely on an alignment of the needles comprising the display: no sensation of depth was necessary, and was indeed absent for two subjects. Furthermore, the stimuli used in all of these early studies of motion parallax were comparatively sparse in terms of the extent of information they provided. Specifically, observers were generally required to make judgements concerning stimuli that contained two single elements, each having a given velocity.

In the 1950s, work of a more interesting and far-reaching nature was carried out by Gibson and his associates. Gibson provided a rich description of how an observer can use motion to gain information about the spatial layout of their environment:

"... as an observer moves through the environment, the projection of the environment in the optic array undergoes a complex, continuing transformation arising from the continually shifting viewpoint from which it is being viewed..".

These optic array transformations contribute *"... to the perception of the fixed spatial layout of the environment"*. Gibson's description of the use of motion in the perception of spatial layout involved the notion of a motion *gradient*. This is produced when there is relative movement between an observer and a textured surface; Gibson suggested that such gradients of optic flow should be called 'motion perspective', in order to distinguish them from the relative optical motions of two isolated objects in space to which the term motion parallax usually refers. However, the latter term is generally used to refer to both cases.

Initial attempts to create the impression of depth through 'motion perspective' were unsuccessful. Gibson and Carel (1952) investigated the role of velocity flow in simulating a receding plane. A 6 foot disc was

covered with narrow line segments of luminous paint which radiated from its centre. When the disc rotated, the velocity of each element was proportional to its distance from the top of the display. A screen covered the disc, and observers looked through horizontal slits at the bottom of this screen. Despite the existence of the velocity gradient, observers failed to perceive this stimulus as a surface, a fact perhaps attributable to its sparseness. Smith and Smith (1957) investigated percepts of convexity or curvature in a textured surface using a variety of cues, including motion gradients. Although motion contributed to the percept of convexity, it did not cause a surface otherwise judged to be flat to look curved.

Using the shadow projection method, Gibson et al. (1959) investigated the percepts evoked by motion parallax with two velocities, and parallax with a flow of velocities. The former case was examined under two conditions. The motion was carried by either two adjacent spots, one above the other, or by many elements — talcum powder was sprinkled onto the surfaces of the two transparent mounts through which the light was to be shone. The ratio of the slower to the faster velocity was either 0.51 or 0.97. It was found that the large velocity difference always gave rise to a percept of separation in depth for the textured field, but that it did not always give rise to such a percept with the spot field. The smaller velocity field did not produce separation in depth for the textured field, and only 7 of the 26 observers of the spot field reported separation. This was probably as a result of positional cues in the latter stimulus. The depth order judgements were sometimes incorrect. Seven of the 26 observers (27%) of the spot field, and 10 of the 46 observers (22%) of the textured field saw the element(s) having the slower velocity as nearer. Spontaneous reversals were occasionally perceived in the latter stimulus.

Gibson et al. (1959) next investigated the depth percepts evoked by a flow of velocities. By projecting light through a horizontally translating, paint-spattered translucent sheet fixed at 45° from the vertical, a vertical gradient of horizontal velocities was created. Nineteen of 21 observers (90%) reported this as a surface slanting away from them at the top of the field, a percept consistent with the motion gradient. Stationary texture was not perceived as slanted. These experiments demonstrate that relative motion parallax gives rise to more accurate percepts in a situation involving a continuous surface which varies in depth than in a situation

involving two separate surfaces.

In accord with previous studies, neither two-velocity displays nor those containing a gradient of velocities allowed accurate judgement of absolute distance.

Flock (1964) also used the point-source shadow projector technique to display various slanted, translating surfaces. Physical slants of -40° to +40° were used. It was found that judged slant was generally underestimated for the greater angles. The relationship between physical and judged slant was not significantly altered by changes in the velocity or duration of the display, or by the nature of the surface (regular/irregular texture).

Computer animation was introduced into perceptual research by Green (1961) and resulted in considerable increases in the degree of control which could be exercised over experimental stimuli and procedures. Using computer-generated displays of moving spots, Braunstein (1968) provided further evidence concerning the ability of velocity gradients to provide information about the characteristics of 3-D surfaces. When texture gradient information and motion perspective information were independently manipulated and the former gradient specified that the surface was in the frontal plane whilst the latter did not, motion perspective was nevertheless capable of producing judgements of slant.

Farber and McConkie (1979) presented observers with computer-simulated translating dihedral angles, which pointed either towards or away from the observer. It was found that observers' depth order judgements were often inaccurate. Subsequent studies have also found a degree of ambiguity in the perceived depth order of parallax displays viewed by a stationary observer (Braunstein and Andersen, 1981; Braunstein and Tittle, 1988, 1993).

The efficacy of observer-produced parallax remained unclear until a pioneering study by Rogers and Graham (1979). Observers viewed a static random dot pattern on the face of an oscilloscope. When the observer was stationary, the dots remained stationary also, and appeared to lie in the plane of the screen. However, side-to-side head movements of the observer generated relative horizontal dot motion on the screen, creating

a vivid impression of a 3-D surface. The extent of relative dot motion was modulated according to either a sine, square, triangle or sawtooth wave function. A condition in which the observer remained stationary and the oscilloscope moved from side to side also produced these unambiguous depth percepts. This seminal study exemplifies the notion that observer-produced motion parallax can provide depth information *independently* of all other sources. A decomposition of the sources which allow unambiguous perception in this situation was carried out by Rogers and Rogers (1992), and is considered in Chapter 6.

1.3.6 Absolute depth judgements and depth constancy

Despite the fact that motion parallax gives rise to compelling percepts of relative depth, parallax can not specify *scalar* depth information unless additional information can be extracted from the visual scene. Richards (1985) has shown that the depth between two rigid points can be obtained from one orthographic stereo view of the two points and the points' velocities. Rogers and Howard (1995) note that the acceleration component of the flow field (in situations involving more than two frames) can theoretically provide information about the change in viewing direction and can thereby specify the shape of the object, given that the latter subtends an angle greater than about 5°.

An alternative source of information allowing the theoretical computation of scalar depth is the absolute distance to the object or objects in question. Although this information is theoretically available in the absolute motion parallax of a point, it was noted above that observers are not adept at extracting this information (e.g. Park, 1964; Dees, 1966; Gogel, 1969; Erikkson, 1972; Ferris, 1972; Gogel and Tietz, 1973). Therefore in order for viewers to extract scalar depth through absolute distance information, they must use alternative cues to the latter.

It has long been known that stereoscopic scalar depth results from the processing of binocular disparity and absolute distance information (e.g. Wallach and Zuckerman, 1963). The relationship between disparity, stereoscopic depth and distance is given by the inverse square law, the derivation of which was presented in Section 1.3.4. In that section an analogous formula for the absolute depth of parallax surfaces was also

derived.

Ono, Rivest and Ono (1986) examined depth scaling of head-linked motion parallax surfaces. Observers estimated the depth of stimuli presented at 40 and 80 cm. The extent of proximal parallax remained constant over all conditions. It was found that the magnitude of depth perceived at the greater viewing distance was between twice and three times that perceived at the smaller distance, indicating that the visual system does indeed calibrate motion parallax according to absolute distance information. However, since the inverse square law predicts that a doubling of viewing distance leads to a quadrupling in perceived depth, this scaling is obviously imperfect.

A second experiment examined the extent of motion parallax required to produce a depth percept of given magnitude at the same distances as above. Observers were asked to adjust the apparent depth of the surface until it was equal to that of a hand-held block of wood. It was found that the amount of parallax required was considerably lower for the displays at the greater viewing distance. The inverse square law predicts that the amount of parallax required for the display at 80 cm would be a quarter of that required for the display at 40 cm. The actual ratio obtained was slightly greater than this, but was nevertheless considerably below a value of unity, providing further support for the theory that the visual system calibrates motion parallax according to absolute distance information.

In a final experiment, three extents of distal parallax were held constant as viewing distance varied between 40 and 320 cm. Observers perceived the greatest amounts of depth in the surfaces at the furthest viewing distance as predicted, but the magnitude of underestimation increased with viewing distance such that 14 arc min equivalent disparity (see Section 1.3.4) produced perceived depths of 5.2 cm and 17.7 cm at distances of 80 cm and 320 cm respectively. This represents a ratio of approximately 3.5, which is evidently well below 16, the ratio predicted by the inverse square law for a quadrupling of distance.

It was found that increases in either viewing distance or the extent of distal parallax led to increased incidences of perceived rocking motion. This was especially true for larger amounts of parallax, and for the greater

viewing distances. The apparent trade-off between perceived depth and rotation is a point which is discussed later in this thesis, briefly in Section 1.6.2 and at length in Chapter 5.

It was noted above that observers can not gain absolute distance information from absolute parallax. Rivest, Ono and Saida (1989) examined depth scaling when various alternative sources of absolute distance information were made available to the observer. In their first experiment, the parallax display was viewed at 80 cm, but convergence angle was manipulated to simulate distances of 40, 60 and 80 cm. It was found that the amount of depth perceived was relatively independent of the convergence angle, indicating that these manipulations had little effect.

In a second experiment, viewing distance was again 80 cm but conflicting size cues were used to manipulate apparent distance. One portion of the display consisted of the parallax stimulus. The other portion showed a Canadian dollar bill presented at either 100% or 70% of its normal size. It was expected that the former condition would result in near-veridical perception of absolute distance, whilst the latter condition would result in a greater perceived absolute distance. Observers were once more asked to report the magnitude of depth perceived in the parallax stimulus. It was found that perceived depth was greater when the bills were presented at the reduced size than when they were presented at their normal size.

In a final experiment, an induction screen was used to alter the apparent distance between the observer and the display. Observers were required to adjust the amount of relative motion in the display until the perceived depth was equal to that in a standard stimulus. In addition, measurements of the apparent distance of the display were taken in order to ascertain the effectiveness of the induction screen. It was found that more parallax was required to produce a percept of a given depth magnitude when the display was viewed through the induction screen than when it was viewed alone. Measurements of the perceived distance of the display indicated that it did indeed appear closer when viewed through the induction screen.

From the results of these experiments it was concluded that motion parallax is calibrated with absolute distance. The lack of depth scaling observed when convergence simulated various viewing distances was due to the inability of this cue to affect perceived distance. The manipulation of perceived distance through size scaling and use of the induction screen illustrate that depth scaling occurs in situations in which this factor has been modified successfully.

In sum, these studies suggest that depth scaling of motion parallax surfaces is fairly accurate up to distances of approximately 1 m but that performance (using artificial stimuli) beyond this distance is poor, perhaps due to the increasing amounts of 'rocking' or rotation of the surface which is possibly perceived at the expense of depth. Unsurprisingly, the scaling is reliant upon cues which influence estimates of the viewing distance; it remains unaffected by cues which merely change the simulated, rather than perceived, distance (Rivest at al., 1989).

Comparisons of judged depth using simulated and real objects has found that judgements are more accurate in the latter case. Durgin, Proffit, Reinke and Olson (1995) firstly presented observers with cones of various simulated tip-to-base depths at a viewing distance of 72 cm. Observers manipulated a mouse-adjustable icon to match the perceived cone profile. A bias was found such that the cones were generally perceived to be approximately as deep as the width of their base. This accords with an earlier suggestion that, in the absence of sufficient information, objects are generally perceived to be about as deep as they are wide — the so-called 'compactness assumption' (Caudek and Proffitt, 1993). When real cones were viewed from a distance of 2 m in a brightly lit, fully structured environment, performance was much improved. Durgin et al. (1995) suggested that this was probably a consequence of the active head motion of the observer which, along with the optic flow from the surroundings, 'reduced the ambiguity of the object's actual lack of rotational motion' (p. 686). Additionally, the change in viewing angle was greater in the latter condition. It must, however, be remembered that the use of real objects does not always *isolate* the information provided by motion parallax in the way that artificial stimuli are able to.

Durgin et al. (1995) also examined depth judgements under conditions

of binocular, stationary viewing. Performance was near-perfect in the conditions comparable to those above. A further experiment, which used real, binocularly-viewed cones at viewing distances of up to 3 m, demonstrated that stereoscopic depth constancy exists over a significantly greater range of distances than does parallax depth constancy.

This discrepancy between performances in the two viewing conditions contrasts with a variety of similarities between motion parallax and stereopsis. These are discussed below.

1.4 The similarity of motion parallax and stereo

1.4.1 Depth/curvature thresholds

It was earlier noted that the geometries of motion parallax and binocular stereopsis are formally equivalent. This has led to investigations of their similarity as perceptual sources of depth information. Thresholds for perceiving the 3-D structure of sinusoidal corrugations were measured as a function of corrugation frequency by Rogers and Graham (1982). Parallax thresholds were expressed in terms of 'equivalent disparity' (see Section 1.3.4). Sensitivity functions for stereoscopic and motion parallax corrugations were similar. Peak sensitivity lay at approximately 0.4 cycles per degree (cpd), with a fall-off at higher and lower frequencies. Rogers and Graham (1985) showed that the fall-off at lower frequencies was not due to the smaller number of cycles on the screen: sensitivity decreased for surfaces having a single Difference-of-Gaussians (DOG) profile which extended over a large spatial area. Motion parallax thresholds were typically double those of stereo. The lowest thresholds obtained in this study were approximately 40 arc sec equivalent disparity. Employing a forced-choice procedure instead of the method of ascending limits used by Rogers and Graham (1982), Bradshaw and Rogers (1993) obtained parallax thresholds as low as 8-10 arc sec equivalent disparity, and stereo thresholds of 2-4 arc sec. Prince and Rogers (1998) examined stereoscopic depth thresholds as a function of eccentricity. They found that as eccentricity increased from 3.5° - 21° the peak-to-trough thresholds increased. The optimal corrugation frequency for detection decreased and the upper cutoff frequency decreased. When the functions were M-scaled, i.e. re-scaled for

cortical size (Rovamo and Virsu, 1979), the peak detection frequency remained constant with eccentricity at a value of 0.8 cycles per cortical mm. It was suggested that the shifts in the disparity modulation sensitivity as a function of eccentricity can be accounted for by the cortical magnification factor.

Johns (1998) has shown that sensitivity to moving cyclopean forms exhibits a similar corrugation sensitivity function to those of parallax and stereo, peaking between 0.4 and 0.8 cpd. Swash and Rogers (1998) showed that the lower limits of transparency for stereo and motion parallax surfaces were ~ 30 arc sec and ~ 2 arc min equivalent disparity respectively. The upper limit for stereo was approximately 80 arc min, while the upper limit for motion parallax exceeded 120 arc min.

Cornilleau-Pérès and Droulez (1993) investigated whether curvature discrimination was better for parallax- or disparity-defined surfaces. Parallax surfaces underwent a rotation of ±12.45° about the vertical axis. Observers indicated whether the presented surface was curved or planar. It was found that thresholds were lower for motion parallax than for stereo, and were lower still for surfaces which contained both cues. This pattern of results is opposite to that obtained by Rogers and Graham (1982), and the authors suggested that a possible reason for the unexpectedly high stereo thresholds was the lower dot density of their displays. A further possible factor is the temporal interpolation procedure used to create disparities smaller than the 2 arc min pixel size in contrast with the theoretically infinite resolution of Rogers and Graham's (1982) analogue displays (Howard and Rogers, 1995).

Further evidence of similarities between thresholds of motion parallax and disparity is provided by experiments involving shape index discriminations. Koenderink (1990) classified smooth surfaces according to the shape index and the curvedness of local surface patches. De Vries, Kappers and Koenderink (1993) demonstrated that observers can classify the shape index of a surface independently of the curvedness. Van Damme, Oosterhoff and van der Grind (1994) examined the ability of observers to classify and discriminate between various surface defined through head-linked motion parallax. Shape discrimination thresholds were lowest for cylindrical surfaces of shape index ±0.5. Thresholds were higher for 'saddle' shapes. These were also the hardest surfaces to classify when defined by disparity (de Vries et al., 1993; de Vries, Kappers and

Koenderink (1994). Just Noticeable Differences (JNDs) were lowest for cylinders, with values of approximately 0.025 and 0.015 for the parallax- and disparity-defined shapes, respectively. These results parallel the depth threshold data obtained by Rogers and Graham (1982), which showed parallax thresholds to be higher than stereo thresholds by a factor of approximately two.

1.4.2 Subthreshold summation

Sub-threshold summation of motion parallax and stereo was investigated by Bradshaw and Rogers (1996). Thresholds for discriminating the phase of sinusoidal corrugations of frequency 0.1, 0.2 or 0.4 cpd were determined for each cue alone. Thresholds were then determined for surfaces containing both cues. The amplitude of each cue was scaled according to its individual threshold. It was found that thresholds for the compound surfaces were approximately a factor of two lower than the normalised thresholds for the separate cue surfaces.

1.4.3 Between-cue threshold elevation

Prolonged viewing of a 3-D surface, the image of which is stabilised on the retina, leads to depth aftereffects. Schumer and Ganz (1979) presented observers with horizontal disparity corrugations defined by dynamic visual noise. The corrugation frequency was either 0.52 or 1.57 cpd. Observers viewed the display for 5 minutes, and were instructed to more their eyes up and down the corrugations. No depth aftereffects were seen on subsequent viewing of a flat test surface, because any given retinal region had been stimulated by both positive and negative disparity gradients, and crossed and uncrossed disparities. However, the depth thresholds of corrugation frequencies identical to the adapting frequencies were found to be elevated by up to 60%. The full bandwidth of the effect at half amplitude was 2-3 octaves. These results were interpreted as indicating the existence of broadly-tuned disparity-corrugation channels, with each channel maximally sensitive to a given frequency. Bradshaw and Rogers (1993) investigated whether stereo thresholds could be elevated through inspection of a parallax surface, and vice versa. Thresholds were firstly determined for a 0.2 cpd surface defined by each cue alone. Observers then adapted to a 4.5 arc min disparity-defined

corrugation, which underwent a phase reversal every 2 seconds, for 3 minutes. Stereo thresholds for the same corrugation frequency were subsequently found to have risen by an average of 112%. When the adaptation surface was parallax-defined, subsequently-measured parallax thresholds had risen by 50%. When the adapting surface was disparity-defined, parallax thresholds rose by 45%, demonstrating that between-cue threshold elevation had occurred. Similarly, a parallax-defined adapting surface led to increases of 50% in the disparity threshold. The authors concluded that there must be a quantitative interaction between the disparity and parallax mechanisms at a relatively early stage in the depth extraction process.

1.4.4 Induced effect

When a surface in the frontal plane has its horizontal size increased in one of the eyes' images, it appears slanted about the vertical axis. The surface appears to slant away from the eye receiving the smaller image. Ogle (1938) termed this the geometric effect, since it is predicted by a consideration of the relative sizes of the images in the two eyes of a real, slanted plane.

If a surface in the frontal plane has its vertical size increased in one of the eyes' images, it again appears slanted about the vertical axis. The surface appears to slant away from the eye receiving the *larger* image (Lippincott, 1889; Green 1889; Ogle, 1938). This is known as the induced effect (Ogle, 1938). The fact that observers perceive the surface as slanted in the latter case, despite the absence of horizontal disparities, led Mayhew and Longuet-Higgins (1982) to suggest that the visual system utilises vertical disparity to provide information concerning absolute distance. The same perceptual effect has also been demonstrated with motion parallax surfaces (Rogers and Koenderink, 1986). Observers viewed a frontal random dot pattern while making side-to-side head movements. The vertical extent of the image increased as the head moved to the right, and decreased as the head moved to the left. The right-hand edge of the surface was perceived to lie closer than the left-hand edge. This percept was reversed when the magnification occurred as the head moved to the left. In a stereo display, one eye received the magnified image and the other received the minified image; the percept was of a surface slanted in

depth, with the nearer side corresponding to the eye receiving the larger image.

There are at least two possible explanations of this result. The binocular vertical size differences or the vertical size changes occurring over time could theoretically be used to calculate the angle of eccentric gaze (Mayhew and Longuet-Higgins, 1982). Alternatively, the visual system may compute surface slant through the extraction of the deformation component of the disparity and flow fields. Decomposition of the vertical magnification into the differential invariants (Koenderink and van Doorn, 1976) leads to dilatation and deformation components. Perceived slant in both the disparity and the motion parallax cases is compatible with the use of the deformation component by the visual system, whilst the dilatation component could indicate the eccentricity of the surface. It was reported by Rogers and Koenderink (1986) that the parallax surface seemed to approach and recede with the side-to-side head movements. This has led to the suggestion that an alternative use of the dilatation component could be to signal the changes in viewing distance occurring in synchrony with the head motion (Howard and Rogers, 1995).

1.4.5 Contrast effects

Further evidence for the similarities between parallax and stereo depth mechanisms is provided through contrast effects. Graham and Rogers (1982a) found that protracted viewing of a sinusoidally-corrugated surface led to a depth aftereffect — a flat test surface was perceived to contain corrugations of opposite phase to those of the adapting surface (see also Long and Over, 1973). This was the case for surfaces defined both through motion parallax and through binocular disparity. The strength of the aftereffect was measured through a nulling technique: increasing amounts of either parallax or disparity were added to the test surface until it was perceived to be flat. Graham and Rogers (1982b) tested whether the depth aftereffect created by viewing of a parallax-defined surface could be transferred to a disparity-defined surface, and vice versa. In the former condition, observers firstly viewed a motion parallax-defined corrugated surface of unambiguous depth order for a period of 15 seconds. During the test phase they viewed a stereo-defined surface in which the corrugations were ambiguous. The ambiguity was created through repeating dot

patterns in the horizontal rows, which meant that the dots could be matched either in front of or behind the fixation spot. It was found that following the period of unambiguous adaptation, the perceived phase of the stereo corrugations was opposite to that of the parallax corrugations. In the other condition, in which observers firstly adapted to unambiguous stereo-defined corrugations and then regarded an ambiguous parallax surface (created through presenting the horizontally shearing motion in the absence of head movements), the perceived phase of the parallax surface was opposite to that of the adapting stereo surface. It was found that the amount of disparity required to null an aftereffect produced by a parallax-defined surface was less than that required to null an aftereffect produced by a disparity surface. Similarly, the amount of parallax required to null an aftereffect produced by a disparity-defined surface was less than that required to null an aftereffect produced by a parallax surface. A greater amount of parallax was required to null an aftereffect produced by a disparity-defined surface than vice versa.

Simultaneous contrast effects in the two domains have also been investigated. Graham and Rogers (1982) showed that when a frontal surface was flanked by surfaces which either slanted away from the observer at the top ('ground' planes) or the bottom ('sky' planes), it appeared to slant in the opposite direction. This was the case for both disparity- and motion parallax-defined surfaces.

An additional simultaneous contrast effect is the disparity analogue of the Craik-O'Brien-Cornsweet illusion (Anstis, Howard and Rogers, 1978). The illusion is created in the luminance domain when two equiluminant regions are separated by a single sharp peak and a trough in the profile, with smooth luminance gradients extending towards the flanking regions. Observers generally perceive the side adjacent to the light side of the discontinuity as brighter than the side adjacent to the dark side. In the stereo analogue the flanking regions, which are of equal disparity, are perceived to be at different depths (Anstis et al., 1978). Rogers and Graham (1983) compared the size of the illusion in the stereo and motion parallax domains. Additionally, they investigated the dependency of the illusion upon the orientation of the depth discontinuity. The depth discontinuity was 8 arc min. It was found that the size of the illusion was about 35% (i.e. the perceived depth difference between the equidistant flanking regions

was almost 3 arc min disparity) when the discontinuity was vertically-oriented but was less than 5% when the discontinuity was horizontally-oriented. This was true for both parallax and disparity-defined surfaces.

The anisotropy between vertically- and horizontally-oriented surfaces also manifests itself at threshold levels of disparity or relative motion. The depth thresholds of surfaces of corrugation frequency lower than 0.1 cpd are greater for vertical than for horizontal sinusoidal depth corrugations (Rogers and Graham, 1983).

These studies demonstrate firstly that the contrast effects (both simultaneous and successive) produced by each cue are similar, indicating possible similarities in the processing of the two cues (Anstis et al., 1978; Graham and Rogers, 1982a; Rogers and Graham, 1984). Secondly, the finding that an aftereffect in one domain can be nulled by a signal from the other (Graham and Rogers, 1982b) provides additional evidence for a depth representation common to both motion parallax and stereo.

Physiological evidence suggests the existence of cortical cells which are sensitive to both disparity and motion (Poggio and Talbot, 1981; Maunsell and van Essen, 1983). A number of cells in MT and MST respond to particular components of the flow field, for example deformations, rotations and dilations (Tanaka, Fukada and Saito, 1989; Orban, Lagae, Verri, Raiguel, Xiao, Maes and Torre, 1992; Lagae, Maes, Raiguel, Xiao and Orban 1994). The responses of a subset of these are also modulated by disparity. For instance, Roy, Komatsu and Wurtz (1992) found that the directional selectivity of cells in MST was dependent upon whether the motion was presented with crossed or uncrossed disparity. Furthermore, Bradley, Qian and Anderson (1995) showed that MT cells, which are usually directionally selective, respond to motion in the opposite direction to that preferred if the motions are separated by disparity. These studies add to the psychophysical results above, and provide further evidence that depth may have a common representation at some point in the visual pathway.

1.5 The ambiguity of motion parallax

The above discussion highlights the similarities and the interactions between the depth percepts evoked by motion parallax and stereopsis, and is highly suggestive of depth processing mechanisms common to both. However, as previously noted, an important difference between motion parallax and binocular stereopsis is the fact that the former involves spatial changes in the optic array over time, while the latter does not. This has important consequences for the inferences which may be drawn from the two different types of information.

1.5.1 The Rigidity Assumption

When viewing is stationary and binocular, the fact that the two eyes receive a different view of the scene at the same time means that differences in the views can only have arisen from the scene's 3-D structure. When the observer (or the object) translates, the visual system receives successive views of the scene. The changes occurring in the pattern on the projection surface could be attributed to the changes in the observer's viewpoint of a stationary, rigid structure; alternatively they could be interpreted merely as a 2-D array of points undergoing various stretching transformations (Koenderink, 1986; Adelson, 1985). It has been suggested that wherever possible, human observers generally assume that the pattern of relative motion arises from a rigid, 3-D structure. This notion was originally termed the principle of minimum object change (Jansson and Johansson, 1973) but is now more commonly called the rigidity constraint (Johansson, 1977; Ullman, 1979). Application of the constraint to the whole flow field would be foolish however, because the world obviously contains many instances of non-rigid objects. To deal with these, notions involving less extreme rigidity assumptions have been made. These include the "local rigidity hypothesis" (Koenderink, 1986), in which it is assumed that the optic transformations conserve mutual distances measured along the surface, but do not necessarily conserve global distances, and the "piecewise rigidity hypothesis" — pieces of the object remain rigid, but other pieces connecting these may undergo bendings and stretchings (Todd, 1983). These local rigidity hypotheses are perfectly illustrated by a slow-motion sequence of a running cheetah, for instance. Distances between the various

joints of the limbs remain constant and the head remains rigid in its structure, whilst the animal as a whole seems to move with a non-rigid fluidity.

1.5.2 Classification of non-rigid motions

Williams (1993) provided a useful categorisation of the various types of non-rigidity. He organised non-rigid motions into three classes: stretching deformations, locally-rigid bending and fragmentation (or a lack of 'coherence').

Stretching deformations were divided into two subsets, superficial deformation and structural deformation. Non-rigidity of the former type manifests itself as fluidly moving surface features on the surface of a 3-D structure of constant form, e.g. 'ants crawling over the surface of an object' (Braunstein and Andersen, 1984). The latter type of non-rigidity, structural deformation, involves a structure which actually changes over time, for example an inflating balloon.

Locally rigid bending refers to instances in which local distances are preserved but the overall shape of the object may change. Instances of this type of non-rigidity include the opening or closing of a paper fan and the previous example of the running cheetah. An algorithm for solving SFM problems in the face of such non-rigidity is provided by Koenderink and van Doorn (1986).

The third class of non-rigidity, fragmentation, refers to cases in which parts of the scene appear to be separated from each other. It is possible to treat as the scene as a single, extremely non-rigid object, but an alternative interpretation exists in which the scene may be interpreted as arising from a greater number of smaller, more rigid structures. Examples of non-rigidity of this nature include shattering glass or splashed water (Williams, 1993).

1.5.3 The perception of non-rigid structures

The frequent occurrence of non-rigidity in the natural environment and our consequent familiarity with such motions should make it

unsurprising that artificial stimuli with no rigid, 3-D solution are indeed easily perceived as non-rigid, 3-D objects. For instance, Jansson and Runeson (1977) presented observers with a wireframe stimulus in which two opposite corners remained stationary whilst the remaining two varied in their position, thus simulating the rotation of a square about its diagonal axis. When a phase lag was introduced between the two moving corners, observers generally perceived the stimulus as undergoing a bending deformation. In the light of the fact that alternative interpretations of the structure were also possible, including ones involving stretching motions, the occurrence of the observed percept lends support to the notion that the visual system prefers solutions which are as rigid as possible. As noted by the authors, bending deformations largely preserve the distances between elements on the structure's surface, implying local rigidity. A host of additional experiments testify to the fact that observers easily perceive non-rigidity when it has been produced in a variety of ways. For example, observers can perceive the bending and stretching of a line of points (Janson, 1977), can perceive various non-rigid human motions (jogging, dancing or walking) when presented with the relative motion of dots whose stationary configuration is ambiguous (e.g. Cutting, Proffitt and Kozlowski, 1978; Kozlowski and Cutting, 1977; Johansson, 1973, 1975, 1976; Todd, 1983) and can perceive the expressions of a human face (Bassili, 1978). Von Fieandt and Gibson (1959) presented subjects with a shadow projected image of a lattice of elastic threads, rotating about its central vertical axis. The stimulus was perceived as a rigid structure. When the vertical perspective component of the projection was removed and the projection consisted only of the horizontal perspective component (i.e. expansion and contraction in the horizontal dimension), the structure was reported as being considerably less rigid.

The perception of non-rigidity need not necessarily impair the observer's ability to accurately perceive form. Todd (1984) asked subjects to discriminate between cylindrical surfaces of various curvatures. These surfaces were presented under both polar and parallel projection, and underwent both rigid and non-rigid transformations. The former transformation consisted of a rotation about the vertical axis; the latter involved a stretching along this axis, in combination with a rotation. It was found that the accuracy of judgements was independent of the rigidity

of the transformation. This suggests that the characteristics of a structure are not necessarily lost when that structure is perceived as deforming. Intuitively this is unsurprising, given that these deformations occurred along the central axis of the structure and did *not* involve variations in the simulated curvature.

The assumption that the visual system chooses to interpret a viewed scene as at least locally rigid gives rise to two questions of interest. The first concerns the occasional tendency of the visual system to interpret a stimulus with a perfectly rigid solution as non-rigid. The second concerns the willingness with which the visual system will entertain a rigid percept when faced with a stimulus which in fact has no rigid solution.

1.5.4 The perception of non-rigidity in rigid structures

A class of stimuli which sometimes result in the former situation are 2-D geometric patterns which, when rotated about the Z axis, may appear as 3-D, non-rigid structures. This phenomenon is known as the stereokinetic effect (SKE), and was probably observed originally by Musatti (1924), and more recently by Wallach, Weisz and Adams (1956), Fischer (1956), Braunstein (1966) and Braunstein and Andersen (1984) (see Section 1.2.2). In the latter study, observers were presented with a three-looped shape which rotated in the frontal plane. Despite the fact that a rigid, 2-D solution was possible, observers reported that the arcs of the stimulus appeared to separate in depth, and that the whole shape appeared as a 3-D, twisting, bending object. The authors proposed that this perception resulted from stimulus size changes, which stimulated automatic processes specifying approach or recession (Regan and Beverley, 1978; Beverley and Regan, 1979) and which were not affected by a rigidity principle.

A further example of a rigid, 2-D translation giving rise to the percept of non-rigidity, this time also 2-D, is provided by Nakayama and Silverman (1988a). Observers were shown a horizontally translating sine wave the ends of which extended beyond the edges of the screen. When the amplitude of the waveform was high, observers perceived rigid, translatory motion. As the amplitude was decreased, there came a point at which the stimulus was perceived as non-rigid. The amplitudes at which

stimuli of higher spatial frequencies became non-rigid were lower than those at which stimuli of lower spatial frequency became non-rigid. In fact, the horizontal movement of the sine wave was perceived as non-rigid if the angle at the zero crossing was less than 15°. The authors provided an explanation based on the intersection of constraint lines defined by local velocity components (Adelson and Movshon, 1982). If noise is added to this model, the judgement of direction of motion is influenced by the angle of local components in relation to the global velocity. When the constraint lines defined by local components are very similar in orientation (as is the case in stimuli of low spatial frequency or low amplitude), a small perturbation in the position of these lines leads to a large change in the position of their intersection. Nakayama and Silverman (1988a) suggest that 'automatic' processes extract the components of the flow field, and Williams (1993) likens this reasoning to that of Braunstein and Andersen (1984), who also noted that the percept of a non-rigid, 3-D structure may have resulted from the operation of automatic processes on the rotating 2-D figures of their study.

Green (1961) presented subjects with projections of either patterns of dots, unconnected lines or connected lines, which rotated about an axis which was either fixed or constantly changing (causing a 'tumbling' motion). Subjects were asked to rate the coherence of each display using a five-point scale, with coherence defined as the extent to which the display elements appeared to remain constant in their positions with respect to each other. It was found that although many subjects perceived the displays as 3-D, less than half the displays were given the highest coherence rating. This suggests that a display with a rigid, 3-D solution may be perceived as non-rigid.

Sparrow and Stine (1998) showed observers four categories of rigidly rotating eight-vertex geometric forms under parallel projection. The first category contained line drawings which, prior to rotation, appeared to be a Necker cube. However, the depth components (z axis coordinates) of the vertices were randomly determined; this became evident only upon the figure's rotation. The second category contained drawings where the vertices were randomly placed along the x, y and z axes. The remaining categories were vertex-only drawings of the two previous categories (the displays consisted of eight dots). Subjects rated the non-rigidity of the

rotating figures. It was found that greatest non-rigidity was perceived in the first category of stimuli, those which initially appeared as Necker cubes. On the basis of these results and two further experiments, the authors suggested that in the absence of monocular static depth cues (such as those present in the 'Necker cube' category) observers use the rigidity assumption. However, when static depth cues are present they may cause the expectation that the structure should appear a certain way upon rotation. Non-rigid percepts are thus produced when these expectations are violated. The notion that static cues are sometimes used has also been suggested by Braunstein, Hoffman, Shapiro, Andersen and Bennett (1987) and Ganis, Casco and Roncato (1993).

A final example of the dominance of non-rigid percepts over legitimate rigid ones is provided by Ramachandran, Cobb and Rogers-Ramachandran (1988). Subjects were shown a SFM display depicting two cylinders with equal diameters but rotating at different angular velocities. Despite the fact that a rigid interpretation was clearly possible, the observers' percept was that the cylinder with the slower rotation speed had a smaller diameter and was rotating at the same speed as the larger cylinder. This obviously necessitated a degree of non-rigidity. In a case in which the narrower of two concentric cylinders rotated more quickly than the larger cylinder, the percept was of a single cylinder which rotated at a constant angular velocity. The narrow cylinder appeared to bulge outwards at its centre. The authors suggested that velocity gradients are treated by the visual system as direct measurements of distance from the rotation axis, even when this results in non-rigidity. The finding that variations in simulated rotation may be interpreted as variations in depth has subsequently been found by Liter, Braunstein and Hoffman (1992, 1993).

These examples illustrate that non-rigid percepts can arise as a result of rigid 2-D motion, in which case the percept may either be 2-D (Nakayama and Silverman, 1988a, 1988b) or 3-D (e.g. Musatti, 1924; Wallach, Weisz and Adams 1956; Fischer 1956; Braunstein, 1966; Braunstein and Andersen, 1984), depending on the stimulus. Non-rigidity can also arise as a result of rigid 3-D motion, where the resulting percept is of a non-rigid 3-D structure (e.g. Green, 1961; Sparrow and Stine, 1998). Finally, two rigid, 3-D structures may be seen as a single, non-rigid 3-D structure if their

relative spatial configuration allows this (Ramachandran et al., 1988).

1.5.5 The perception of rigidity in non-rigid structures

In a consideration of the willingness with which the visual system will entertain a rigid percept when faced with a stimulus which in fact has no rigid solution, it is useful to consider the fact that the vast majority of the projections used in SFM are only approximations of the retinal projection of a real, 3-D object. The accuracy of the approximation and the nature of the information which is lost are dependent upon both the projection method and the type of structural transformation simulated. For instance, under orthographic projection, all information relating to the distance between the projected structure and the observer is destroyed. Ullman (1983) provides a comprehensive evaluation of the nature of information lost as a result of various projection types.

A cursory mental analysis of, for instance, a rigid, rotating SFM cylinder should lead to the conclusion that deviations in the dot motions, if sufficiently slight, would probably go unnoticed by an observer and would therefore not affect the perceived rigidity of the structure. Hogervorst, Kappers and Koenderink (1996) devised a model which obtained the nearest rigid 3-D Euclidean structure to a given non-rigid 3-D stimulus. Re-analysing the results of an earlier experiment by Norman and Todd (1993), they found that subjects often perceived stimuli which in fact had no rigid solution, as rigid structures. In the original study, observers were presented with SFM stimuli consisting of twelve line elements inside a cube with sides of 10 cm. The task was to indicate perceived rigidity in two different conditions. In one condition, objects stretched in the viewing direction and then rotated about an axis in the projection plane. In the other condition they stretched in a direction perpendicular to that of the rotation axis, and then rotated. It was found that subjects perceived non-rigidity in the latter case but not in the former. Hogervorst et al.'s (1996) re-analysis of observers' rigidity responses found that if there was a rigid, Euclidean solution the projected points of which did not deviate significantly from those of the stimulus, then a rigid structure was perceived. The structure which underwent a simulated stretch along the line of sight in the experiment of Norman and Todd (1993) was found to have a rigid, Euclidean solution which rotated with a

variable angular velocity. Although measurements relating to perceived rotation were not taken by Norman and Todd (1993), a subsequent study by Pollick (1997) took such measurements, and found that observers did indeed perceive maxima and minima in perceived angular velocity as predicted by the model. The structure which, in Norman and Todd's (1993) study appeared non-rigid (namely the structure in which the stretching was at right-angles to the line of sight), was found to have no reasonable rigid Euclidean interpretation and thus it is unsurprising that observers perceived non-rigidity.

The results of other studies examining perceived rigidity often show observers' non-rigidity ratings to be higher than expected. For example, Braunstein (1962) found that objects which were displayed with perspective projection produced increasingly non-rigid percepts as the simulated viewing distance decreased (i.e. as the level of perspective increased). Since the actual viewing distance remained constant, a likely cause of the non-rigidity was the increasing conflict between simulated and actual distances. In terms of the tolerance analysis of Hogervorst et al. (1996), it is likely that the nearest rigid, Euclidean structure became increasingly dissimilar to the actual stimulus, and non-rigidity was perceived as a result.

1.5.6 The validity of the rigidity assumption

The preceding paragraphs have illustrated firstly that the visual system is frequently content to perceive non-rigidity where none necessarily exists, and secondly that the system is content to perceive rigidity when it is in fact absent. However, the various rigidity hypotheses are not invalidated by these observations. In *general,* it is the case that where a rigid solution is possible this solution is usually the preferred one.

1.6 The ambiguity of rigid motions

The ambiguity of optic array transformations which occur over time may not necessarily involve non-rigidity. Many studies have reported that manipulation of a given variable leads to a change in the relative magnitudes of two qualitatively different attributes of a *rigid* surface.

1.6.1 The perceptual interpretation of motion parallax when combined with stereo

Rogers and Collett (1989) presented observers with a horizontally-corrugated test surface which translated to and fro across the observer's line of sight. The corrugations were specified either by motion parallax alone, or by binocular disparity alone, or by the two cues together in varying proportions. The ranges of stereo and motion parallax amplitudes used were 0 - 8 arc min disparity and -11 - +11 arc min equivalent disparity respectively. The observer adjusted the peak-to-trough amplitude in a match surface, which contained congruent amounts of parallax and disparity, until test and match appeared equal in depth. When the disparity and parallax amplitudes were different, the magnitude of perceived depth was dominated by the stereo information, with motion parallax having only a small effect. This parallax information was not ignored however: it was reported that the surface appeared to 'rotate about a vertical axis through the centre of the pattern as it translated to and fro'.

Rogers and Collett (1989) noted that the parallax transformation generated by a real, 3-D surface which translated horizontally in the fronto-parallel plane was approximately consistent with alternative, rigid interpretations. For instance, the same amplitude of relative motion would be generated by a surface containing less depth which rotated about a vertical axis in a 'convex' direction with respect to the observer. The same amplitude could alternatively be generated by a surface with more depth which rotated about a vertical axis in a 'concave' direction with respect to the observer. It could also be generated by a surface of *opposite* depth order the depth of which was greater still and which rotated in a concave direction (See Fig. 1.4).

Observers' percepts of the display were consistent with the geometry shown above. As noted, the perceived depth order was always consistent with the disparity information. When the motion parallax and binocular disparity amplitudes were 'in-phase' (i.e. the gradients of each signal were of equal sign), the parallax information was interpreted as a convex rotation of the surface with respect to the observer. When the amplitudes

were 'out-of-phase' (i.e. the gradients of each signal were of opposite sign), the surface was perceived to rotate in a concave direction.

When the surface was specified by motion parallax alone there was no variation in observers' percepts, despite the variety of legitimate interpretations specified above. A monocularly-viewed motion parallax surface never appeared to rotate but was always perceived to translate in the frontal plane. Its matched depth was consistent with the relative motion amplitude. It seems, therefore, that the variety of possible percepts arising from a given parallax display is realised only when the visual system is constrained by the stereo cue. The visual system is then forced to interpret the parallax information in such a way that it 'fits in' with the stereo information. In the absence of such constraints the system has a bias to interpret the parallax as a translating surface with a specific depth order.

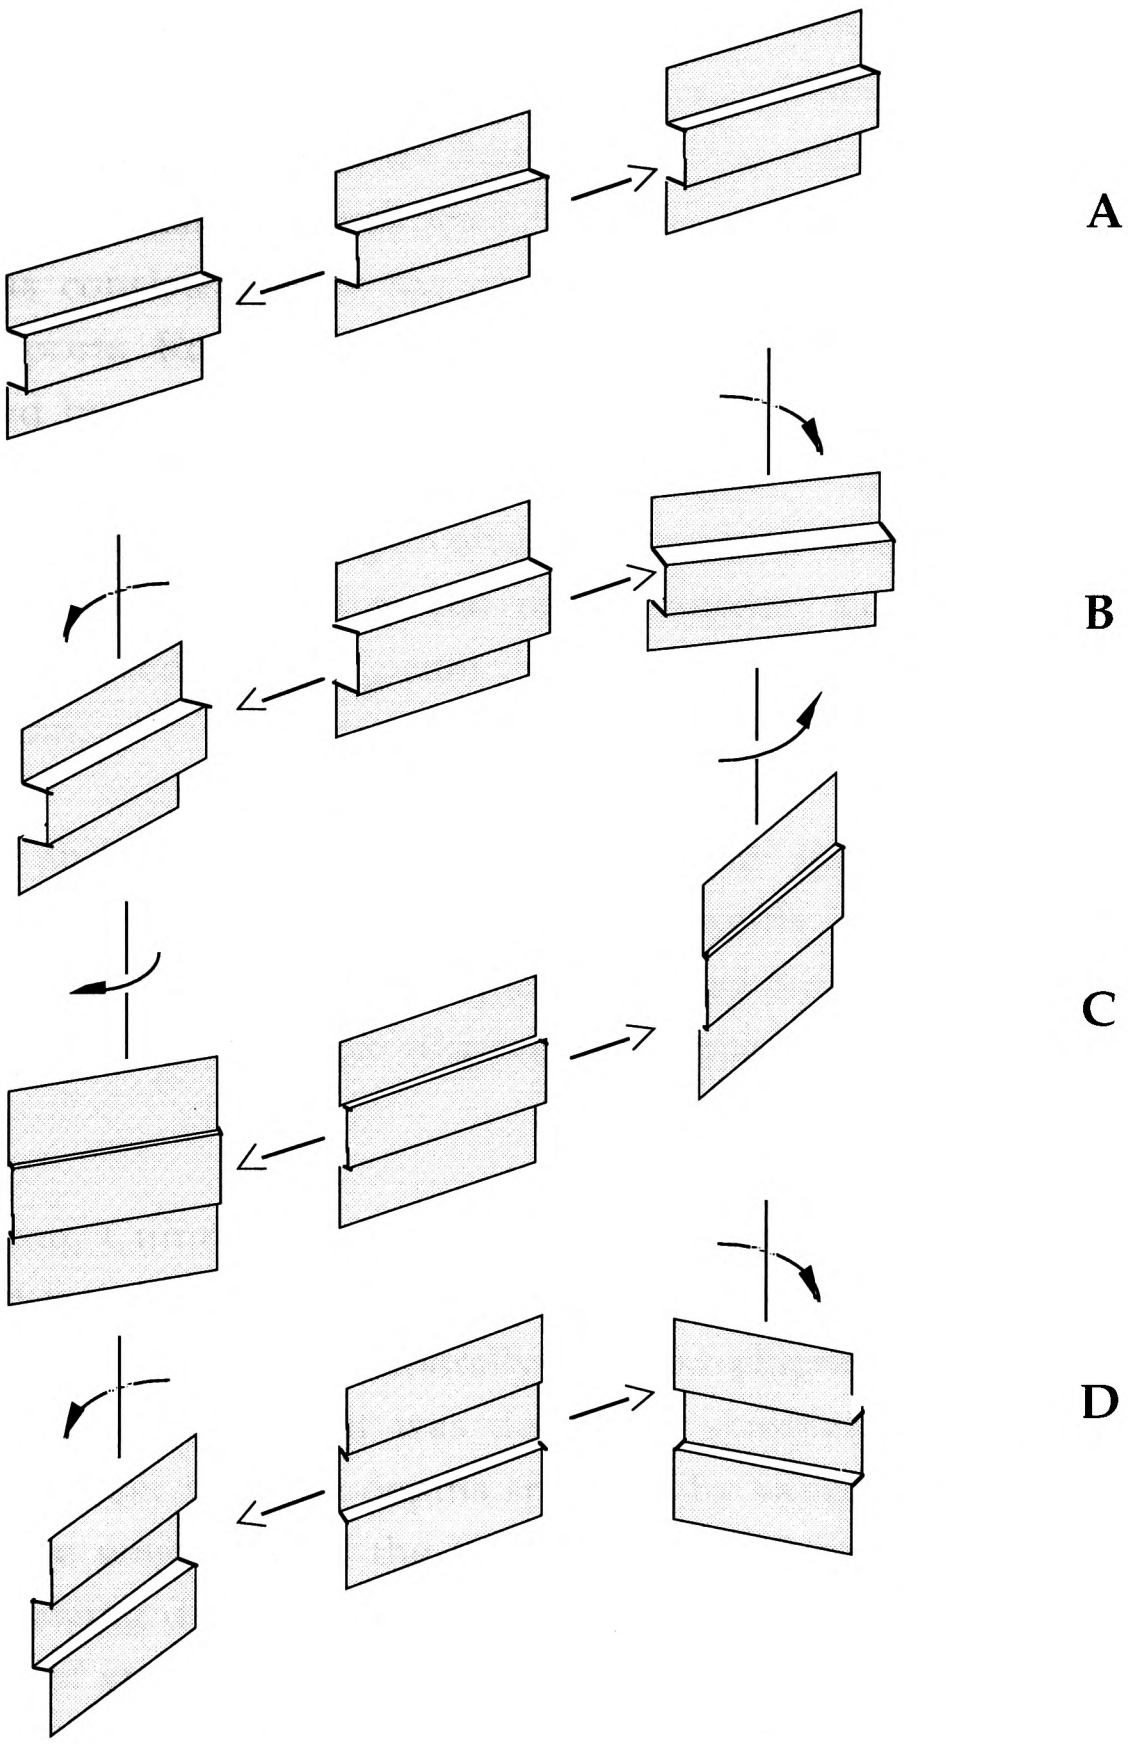


Fig. 1.4: Alternative ways of interpreting a square wave relative motion gradient. To a first approximation, it would be produced by
A: a rigid surface containing a given amount of peak-to-trough depth, which translated in the frontal plane
B: a surface with greater depth than A, which rotated in a concave direction as it translated
C: a surface with less depth than A, which rotated in a convex direction as it translated
D: a surface with reversed depth order, which rotated in a markedly concave direction as it translated.
Adapted from Rogers and Collett (1989, Fig. 3, p. 704).

The greater weight of stereo in the determination of perceived depth order was also observed by Ichikawa and Saida (1996), who used disparity and parallax amplitudes of 0 - 2 and 0 - 20 arc min (equivalent disparity) respectively. When the cues specified different depth orders, perceived order was consistent with the disparity information 80.2% of the time. This supports Rogers and Collett's (1989) conclusion that order is dominated by stereo information even at very low disparity amplitudes. Rotation was perceived over 97% of the time — in a direction against observer movement when perceived depth order was consistent with the parallax information and in the same direction as observer movement when perceived depth order was consistent with the disparity information.

Tittle and Braunstein (1991) also found that stereo information dominated motion. Observers viewed ellipsoids in which the two cues simulated differing extents of depth. Perceived amplitudes were primarily determined by stereoscopic disparities, although motion played a small role also.

Johnston, Cumming and Landy (1994) found that motion and stereo interacted in a weighted linear fashion in a shape judgement task. Subjects were presented with a horizontally-oriented elliptical cylinder, and were asked to judge whether it was deeper or shallower than a cylinder of circular cross section. The depths specified by stereo and motion differed. The relative weighting of the two cues was approximately equal at the nearest viewing distance (50 cm), but at the furthest distance (200 cm) the weighting of stereo was approximately halved; it was thus suggested that the weighted linear combination occurred after the disparities had been correctly scaled for distance.

Tittle and Braunstein (1993) also showed that perceived shape is affected by both stereo and motion. Observers viewed a transparent cylinder under parallel projection, which either rotated about, or translated back and forth along, its horizontal axis. Disparity indicated a depth-to-height ratio of either 1:1, 2:1 or 3:1. Subjects indicated the judged depth-to-height ratio. It was found that judged depth was greater for the rotation displays than for the translation displays, which in turn contained more depth than a static, stereo display.

Norman and Todd (1995) showed that the relative effectiveness of motion and stereo was dependent upon their directions of curvature. Observers were presented with a disparity-defined, horizontally oriented sinusoidal surface. The task was to increase the amplitude of a motion-defined, vertically oriented sinusoidal surface until the surface was visible. In another condition, the disparity- and motion-defined surfaces were vertical and horizontal respectively. It was found that a higher amplitude was required in the former case.

When the observers were presented with a motion-defined surface and increased the amplitude of a disparity-defined surface until it was visible, a greater amplitude was required when the former was horizontal and the latter was vertical than vice versa.

The authors suggested that when faced with conflicting inputs the visual system uses suppression, at least in the conditions of their study. The major factor determining which information was visible and which was suppressed was the relative orientations of the curvatures in each modality.

Further evidence that suppression can occur in certain circumstances is provided by Turner, Braunstein and Andersen (1997). Observers were presented with two intervals. One interval contained a smooth, disparity-defined surface in the presence of motion parallax noise (dots were randomly located in a volume). In the other interval, both cues indicated the volume of randomly located points. The task was to determine the interval in which the surface was present. It was found that performance in this condition was markedly better than in a condition in which the parallax cue indicated the surface and the disparity cue indicated the noise. The authors concluded that when the cues indicated different information about the presence of a smooth surface, the disparity information completely dominated the motion information.

These studies lead to the following tentative conclusions. In situations in which the motion flow field can be resolved with the disparity field in a geometrically legitimate fashion (e.g. Rogers and Collett, 1989; Ichikawa and Saida, 1996), stereo information dominates the perceived depth order when the cues conflict, and is heavily weighted in the magnitude of

perceived depth when the cues are of like sign but different amplitude. It is emphasised that the parallax information in these circumstances was *not* suppressed or 'vetoed' (Bülthoff and Mallot, 1988) — it fitted in with the percept in the form of a rotation. In situations in which the cues specify similar curvatures but different amounts of depth, the visual system combines motion and stereo information in a roughly linear weighted fashion (Johnston et al., 1994; Tittle and Braunstein, 1993). Suppression can occur when there is no rigid solution incorporating the discrepant cues (Turner et al., 1997) or when the orientations of curvature specified by each cue are different (Norman and Todd, 1995).

1.6.2 The perceptual interpretation of motion parallax when presented alone

As noted above, in the Rogers and Collett (1989) study some of the perceptual interpretations of the parallax information came into play only when forced to do so by the constraining stereo information. In the absence of binocular disparity, observers perceived the surface to translate without concomitant rotation, and to have a perceived depth equal to that specified by the relative motion amplitude. Perceived depth order was always the same — dots moving in the same direction as the oscilloscope were perceived as lying on peaks, and dots moving in the opposite direction were perceived as lying on troughs.

In the numerous studies of head-linked motion parallax (e.g. Rogers and Graham, 1979, 1982; Ono and Steinbach, 1990; Ono et al., 1986; Rivest, Ono and Rivest, 1989; Ono and Ujike, 1993; Saida and Ono, 1989) there has not been a single reported instance of the perceived depth order being such that dots moving in the same direction as the head are seen as lying on peaks and dots moving in the opposite direction are seen as lying on troughs. It appears that despite the geometric legitimacy of this solution (with the accompanying concave rotation), the visual system is biased toward solutions containing the depth order noted in the previous paragraph.

There are, however, several reported instances of perceived rotation (e.g. Ono et al., 1986; Ono and Steinbach, 1990; Ono and Ujike, 1993; Saida and Ono, 1989). Studies which have examined this percept have, without

exception, reported it as occurring in a convex direction, that is, against the movement of the observer.

Motion parallax studies involving stationary observers have also demonstrated that rotation is sometimes perceived in displays simulating a translating surface. For instance, Braunstein and Andersen (1981) showed observers a horizontally-oriented dihedral angle formed by two slanted planes meeting at a horizontal line. A representation of the velocity gradients is shown below (Fig. 1.5). The simulated motion of the structure was a horizontal translation, but subjects frequently perceived a rotary component. The authors suggested three cases in which perceived rotation might occur. In the first case, the structure appears reversed in depth, with the slowest dots appearing closest to the observer and the fastest dots appearing furthest away. The velocity gradient is inconsistent with a translation of the surface, but to a first approximation is consistent with a rotation, with the faster dots rotating about a vertical axis through the slower dots. In the second case, the depth order is not reversed but there is an underestimation of surface slant. The authors suggested that this case may be further divided into two subcases. The first consists of percepts arising from stimuli which contain non-stationary dots.

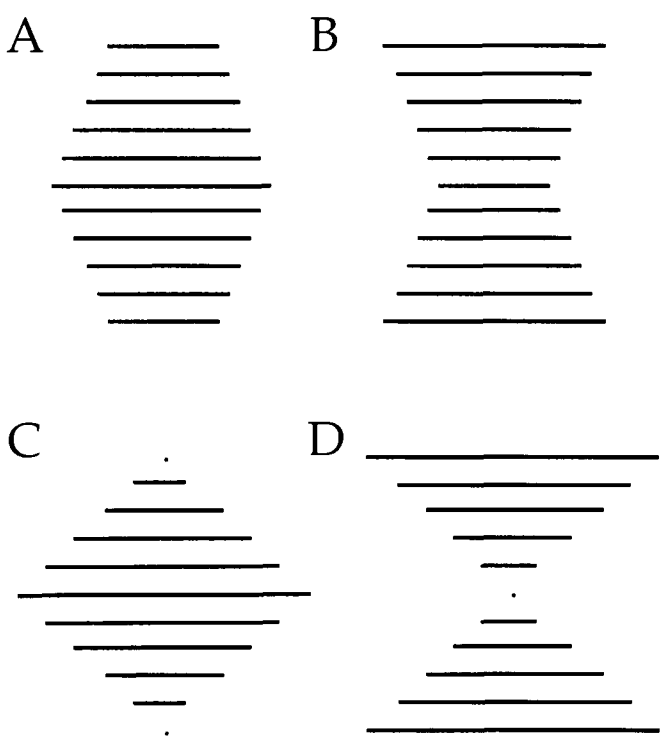


Fig 1.5: Representation of the velocity gradients used by Braunstein and Andersen (1981). Taken from Braunstein and Andersen (1981, Fig. 1, p. 146).

The underestimation of surface slant results in a failure of the translatory motion to account for the difference between the fastest and slowest dots, and thus a rotational component is introduced into the

percept to compensate for the additional portion of the velocity difference. Once again, the faster dots rotate about an axis through the slower dots. The second subcase contains percepts arising from stimuli containing stationary dots. If these are not perceived to lie at an infinite distance then their failure to move is inconsistent with a rigid translation. An alternative percept is a rotation of the faster dots about the slower dots, with no translatory component.

The authors hypothesised that perceived depth reversal might be associated with stronger percepts of rotation. Subjects were instructed to classify the amount of perceived rotation into 'no rotation', 'some rotation' or 'mostly rotation'. Although the proportions of these responses were slightly higher when the structure was seen as reversed in depth, the proportion of high-rotation judgements made with correct depth judgements was sizeable (0.27).

These studies demonstrate the frequent perception of rotation in translating motion parallax displays. Chapters 3 and 4 of this thesis consider the above studies in greater detail and investigate the factors upon which perceived rotation, at both threshold and supra-threshold levels, depends.

Many investigators have noted that the perceived quantities of two qualitatively different attributes often exist in a trade-off relationship, such that an increase in the perceived magnitude of one attribute coincides with a decrease in the perceived magnitude of the other. The principles assumed to underlie these observations are essentially similar to those expounded by Rogers and Collett (1989). Namely, to a first approximation, a given pattern of relative motion may be projected by structures exhibiting various different combinations of depth and global motion. A formal explication of this notion, and an investigation of its implementation by the visual system, is the subject of Chapter 5.

Chapter 6 further examines the putative relationship between the relative magnitudes of a structure's depth and motion, specifically examining the influence of dynamic vertical perspective on these attributes. The final experiment of that chapter investigates the influence of vertical perspective information upon the perceived depth order of

parallax surfaces viewed by a stationary observer.

1.7 Conclusion

The inherently ambiguous nature of the motion parallax cue means that the visual system is theoretically faced with an unlimited number of solutions to a given pattern of retinal image motion. Various sources of information constrain the infinite number of interpretations (within the limits of those defined by the image motion) so heavily that the observer is left with a single, generally stable, percept. This thesis aims to elucidate some of the factors underlying the visual system's biases for one interpretation over another.

Chapter 2:

General Methods

2.1 Hardware

2.1.1 Motion Parallax Display

The apparatus used for the work in this thesis was similar to that used by Rogers and Graham (1979, 1982). A 50% random dot pattern, generated on an Amiga 2000, was displayed on a Hewlett Packard 1304 large screen display oscilloscope. When viewed from 57 cm the display subtended 25° x 20° (horizontal by vertical). Pixel size was 5 arc min x 5 arc min. The random dot pattern was displayed on the oscilloscope screen using a raster scan. The line frequency was 15.6 kHz. The frame frequency was 50 Hz.

This display was viewed monocularly by an observer whose head rested on a chinrest. This chinrest was mounted on castors, and could move laterally through a distance of 13 cm, twice the average ocular distance. In order to simulate the patterns of relative motion that would be produced by a real 3-D surface when the observer translated from side to side, the pattern was systematically distorted with each movement of the observer's head. (Fig 2.1). The distortion was achieved by modulating the amplitude of an additional waveform fed to the x deflection plates of the scope. A potentiometer attached to the chinrest provided the voltage that modulated the distortion waveform. This waveform was generated by a Wavetek 175 arbitrary waveform generator, synchronised to the frame rate of the display. Unless otherwise specified, the waveform was sinusoidal. As the head of the observer moved to the left, several horizontal bands of dots, representing peaks, moved to the right. Bands of dots representing troughs moved to the left. As the head moved to the right, the direction of dot motion was reversed. The direction and extent of motion of the dots lying between the simulated peaks and troughs was interpolated correctly

according to a sinuoidally corrugated surface. The dot pattern was maximally distorted when the observer was at either of the endpoints of the translation path. When the observer was at the exact mid-point of the path, there was zero distortion. When the observer was stationary, no dot motion occurred and no depth could be seen: the display appeared to be 2-D, and to lie in the plane of the oscilloscope screen.

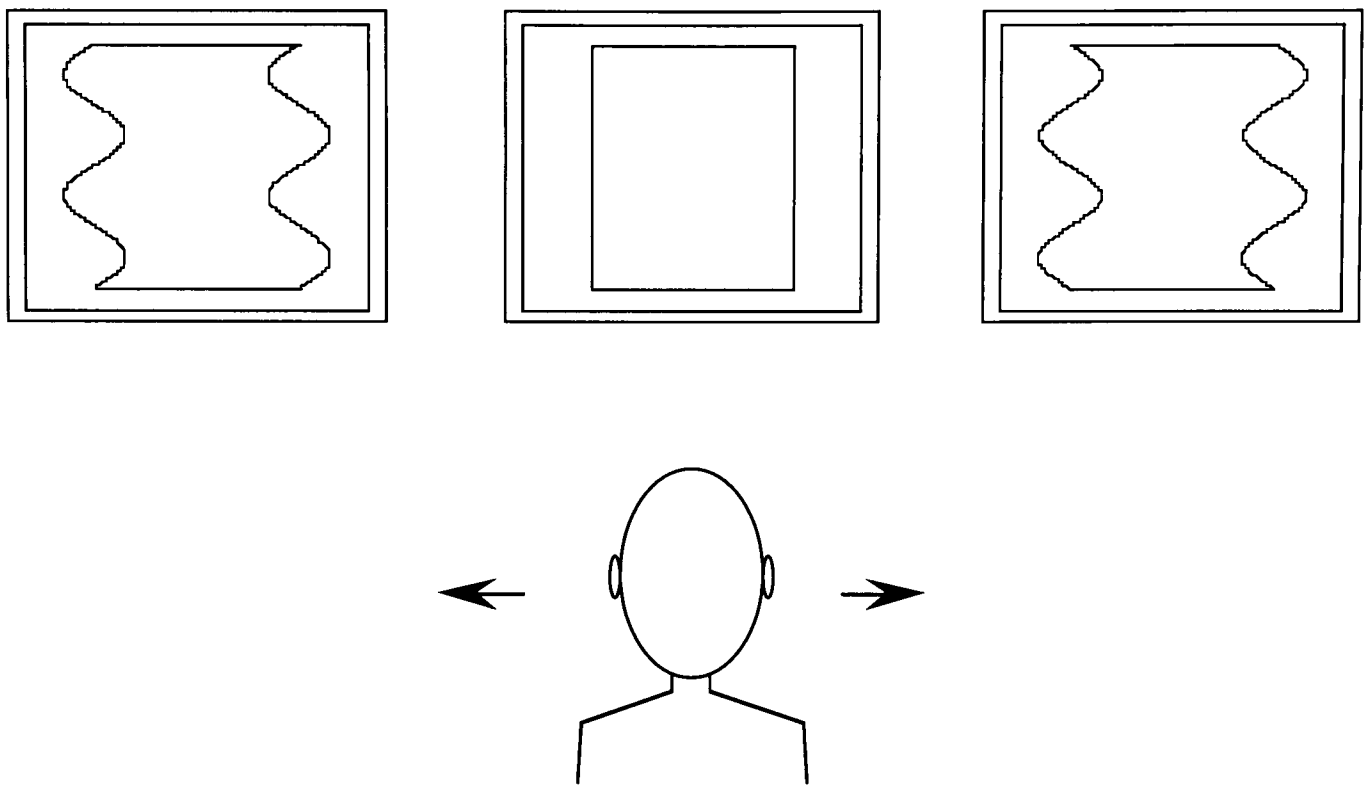


Fig 2.1: Basic experimental set-up, showing the sinusoidal distortions occurring in the display as the observer moved from side to side.

2.1.2 Calibration of Motion Parallax Display

The equipment was calibrated so that a 13 cm translation of the chinrest produced a voltage change from +10.00V to -10.00V from the chinrest potentiometer. This voltage caused any pixel located at either the peak or the trough of the waveform to move a distance of 6 pixels (25 mm) when the amplitude of the waveform was set to 7.50V on the Wavetek 175. When viewing distance was 57.1 cm, these parameters resulted in a relative motion between dots located on peaks and troughs of 4 arc min equivalent disparity per volt. The voltages involved in this and subsequent calibrations were measured by a Fluke 73 Multimeter.

2.1.3 Experimental Sequence Control

A Macintosh II computer controlled the experimental sequence. The desired stimulus parameters were sent to the waveform generator by the Macintosh on each trial. Subjects' responses were recorded by an input/output interface, which had ports for the acquisition of both digital and analogue data. Data from these ports was sent to the Macintosh II, where it was stored for later analysis. Since the interface was generally employed as an Analogue-to-Digital Converter (ADC), it is labelled as such in Fig. 2.2 below, which illustrates the general set-up.

An electronic metronome allowed observers to achieve and maintain a constant head oscillation frequency, which was usually 1 Hz.

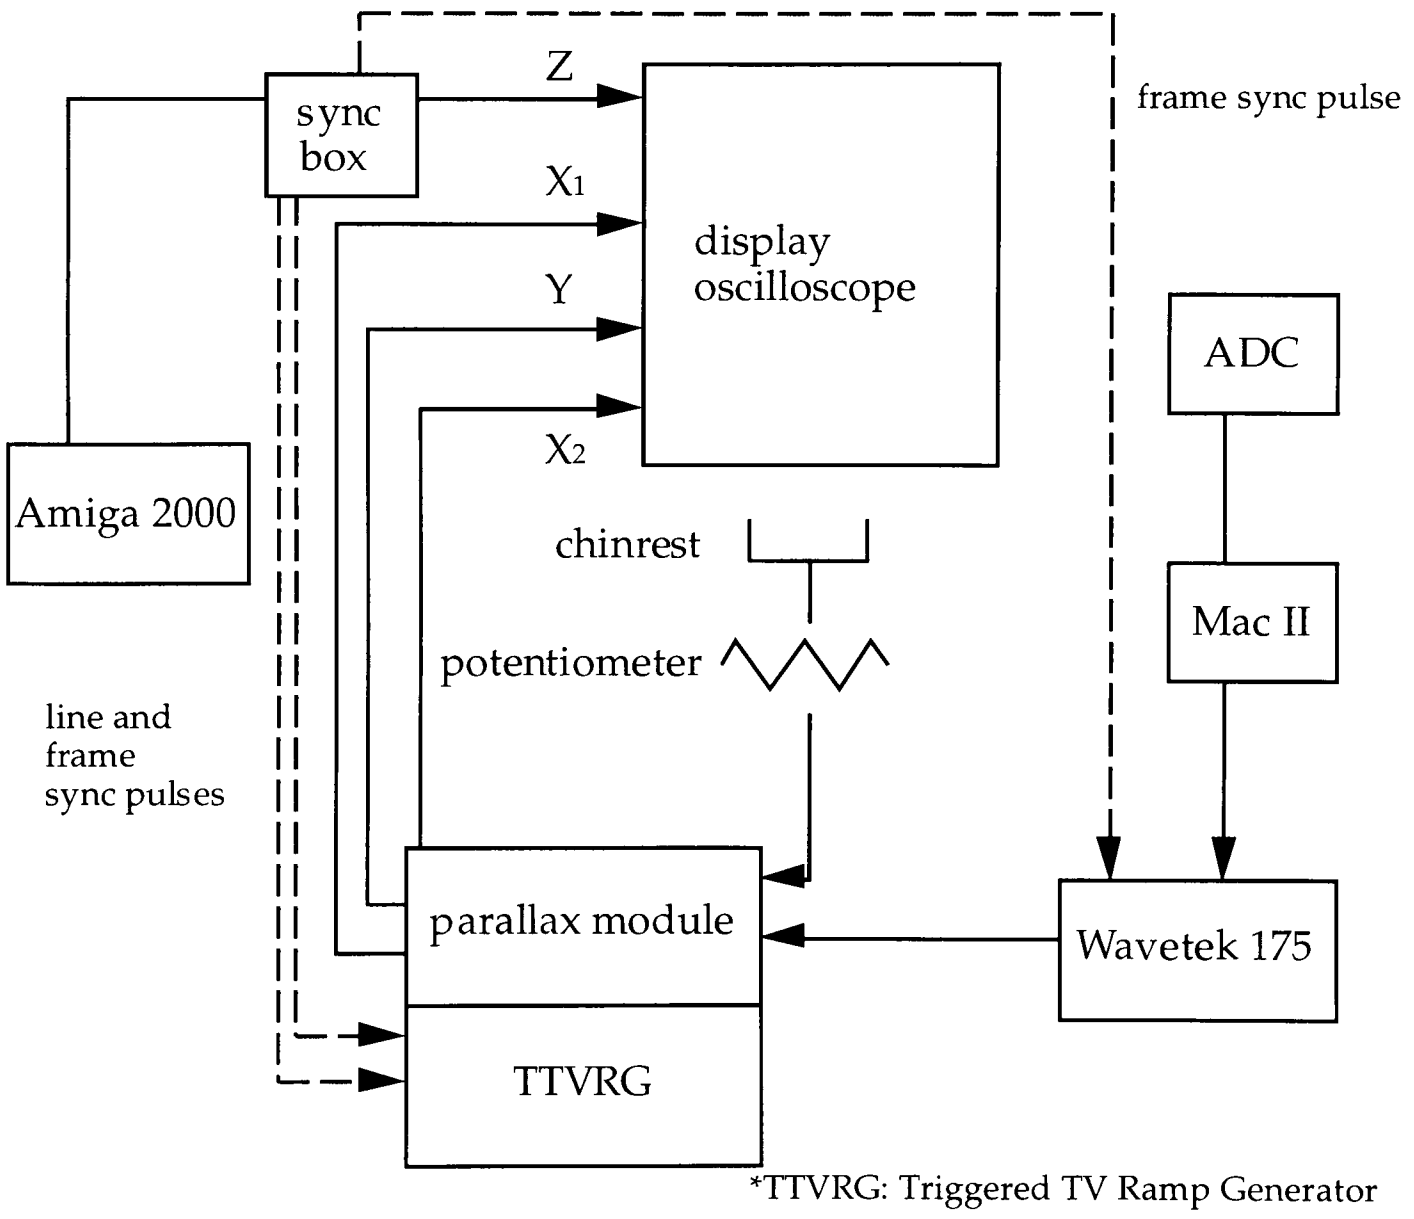


Fig 2.2: Schematic representation of the apparatus used to generate and display the experimental stimuli.

2.1.4 Vertical Image Size Changes

In Chapter 5 of this thesis, dynamic vertical perspective is examined as a cue to rotation. The observer's head movement was accompanied both by motion parallax and by a perspective distortion of the display.

As the eye travelled towards one of the endpoints of its translation path, the vertical size of the display was modulated such that the vertical extent of one side of the display increased whilst that of the opposite side decreased. The vertical size of the intermediate parts of the display was interpolated to create a trapezium. The difference between the sizes of the vertical sides was maximal when the head was at the endpoint of its translation path. Translation of the head to the opposite endpoint resulted in the reverse trapezoidal transformation (Fig. 2.3). These transformations approximate those created by a rotating planar surface. Note that they do not include the horizontal width changes or the texture gradients which occur in the projection of a rotating surface.

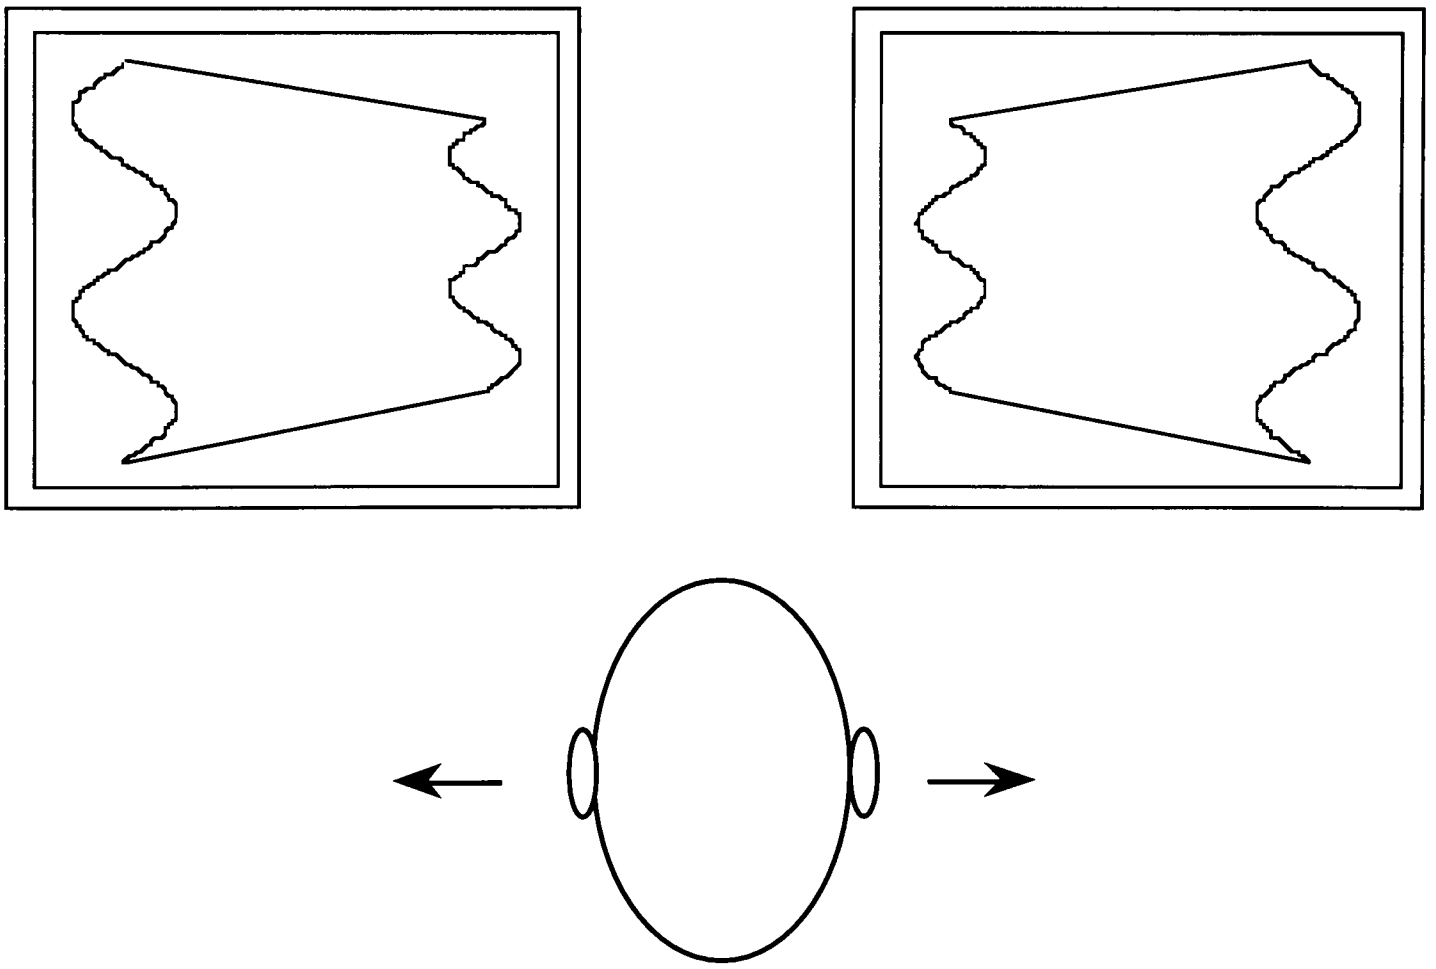


Fig. 2.3: The trapezoidal distortions which occurred with the side-to-side head movements of the observer.

These size changes were accomplished by modulating the amplitude of

the waveform fed to the y deflection plates of the scope. The modulation was carried out according to a ramp waveform which was generated by a Wavetek 186 arbitrary waveform generator. The frequency of the waveform was the same as the line rate of the scan.

2.1.5 Calibration of the Vertical Image Size

In order to vary the simulated rotation angle, the magnitude of the trapezoidal transformation was varied. This was achieved through varying the amplitude of the ramp waveform generated by the Wavetek 186. In order to calibrate the changes in the vertical size of the image due to the amplitude variations, the variation in length of the left- and right-hand sides of a square was measured as the Wavetek 186's voltage was systematically changed.

1. The chinrest was fixed in position such that the voltage across the chin potentiometer was fixed at +10.00V. The amplitude of the ramp waveform generated by the Wavetek 186 was 0.00V.

2. A square of sides 180 mm was displayed on the oscilloscope.

3. The amplitude of the ramp waveform was increased from 0.0 to +2.0V in 0.2V increments. For each voltage, the heights of the left- and right-hand sides of the square were measured.

4. Step 3 was repeated a further three times. Fig. 2.4 shows the calibration data. Each point represents the mean of four separate measurements.

4. Measurements were also made for selected Wavetek 186 amplitudes when the chinrest was fixed in position such that the voltage across the chin potentiometer was fixed at -10.00V. This was in order to ensure that the degree of trapezoidal distortion was symmetrical about the midpoint of the chinrest's translation path.

Fig. 2.5 shows the simulated angle of rotation as a function of voltage. It can be seen that the simulated angle increases linearly with increasing voltage. Appendix A shows the calculation of the required trapezoidal

transformation for a given simulated angle of rotation at a given distance.

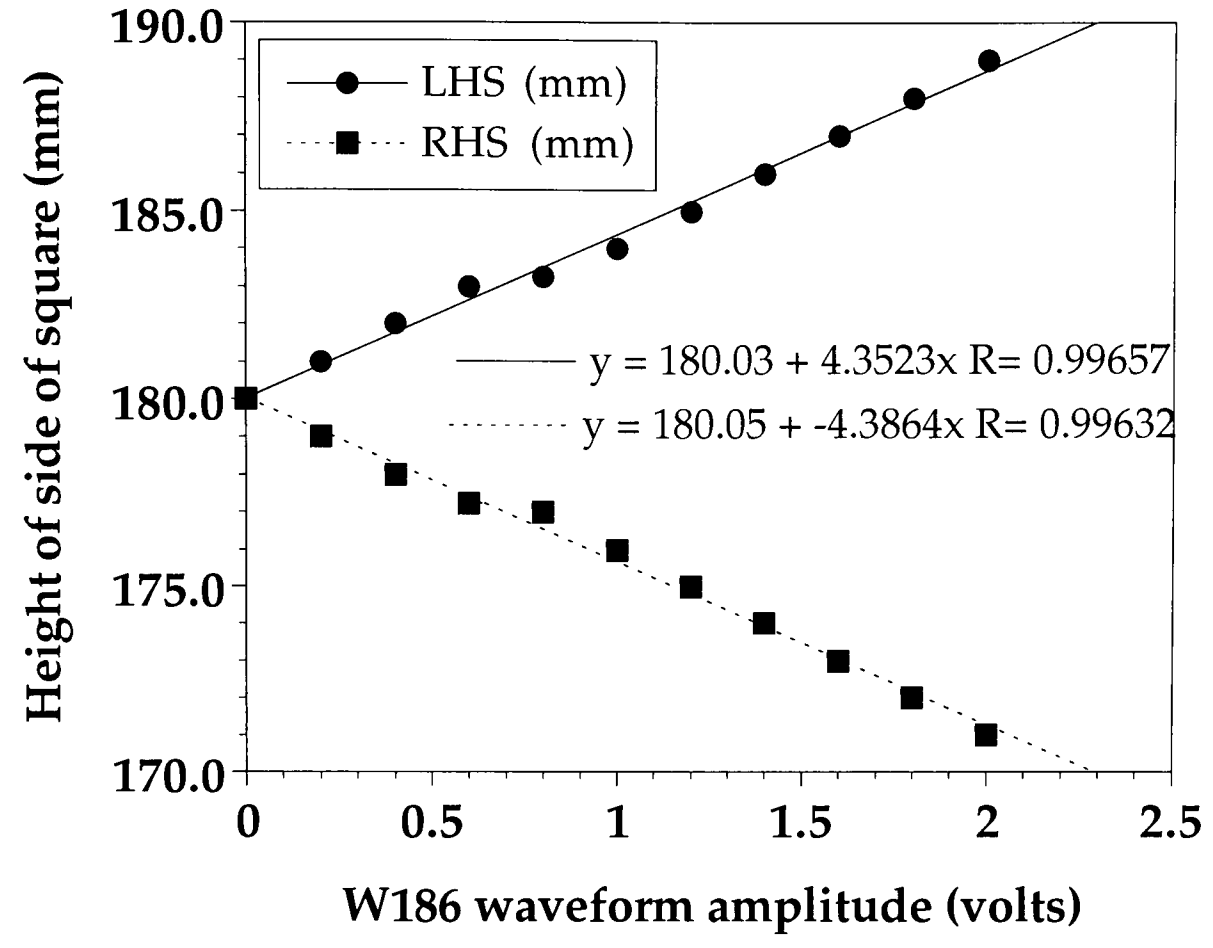


Fig. 2.4: Height of left- and right-hand sides of square as a function of W186 waveform amplitude. The chinrest was held in position such that the voltage across its potentiometer was +10.00V.

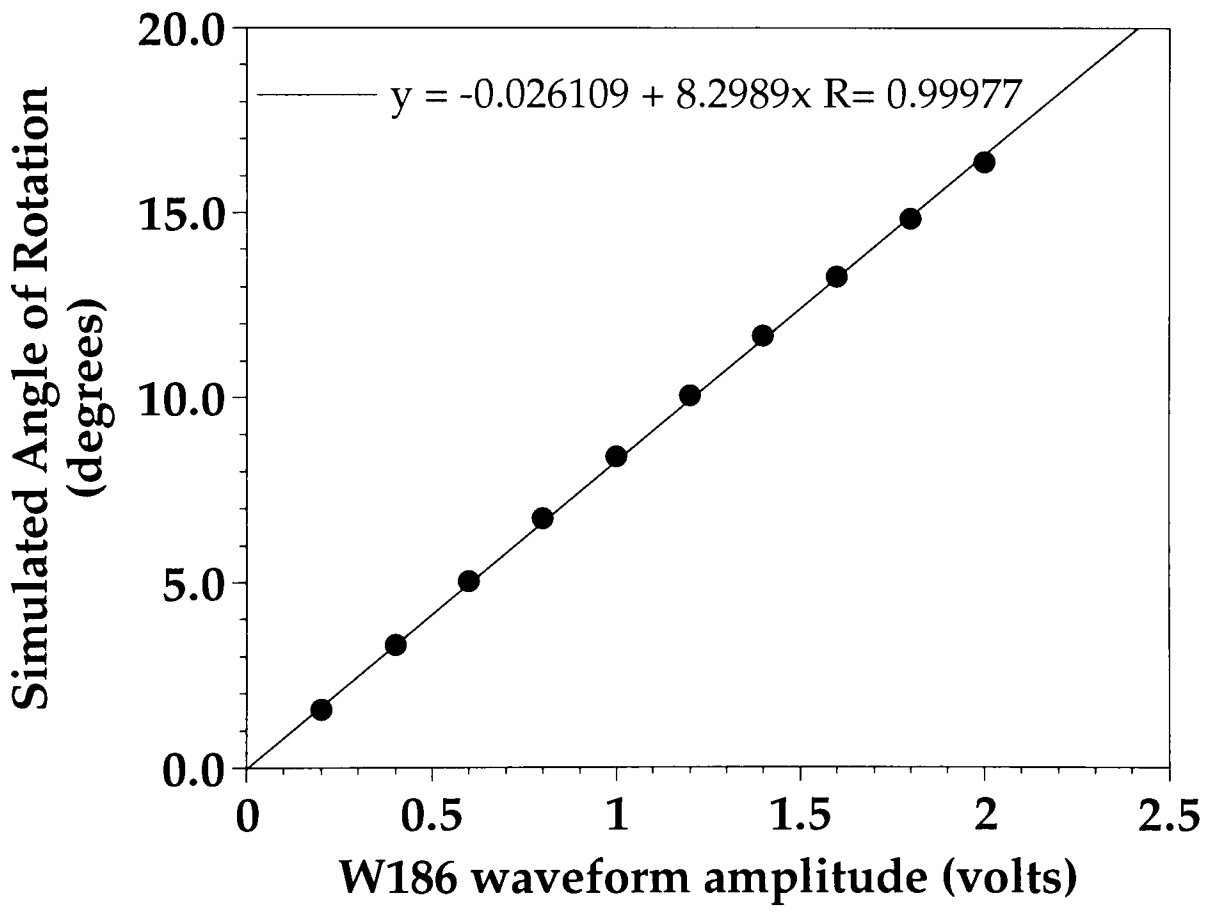


Fig 2.5: Simulated angle of rotation as a function of W186 waveform amplitude when the voltage across the chinrest potentiometer was +10.00V.

In order to check the linearity of the trapezoidal transformation as a function of the chinrest position (measured by the voltage across the chinrest potentiometer), a second calibration was carried out. The variations in length of the left- and right-hand sides of the square were measured as the voltage across the chinrest potentiometer was systematically changed.

1. The amplitude of the W186 waveform was fixed at a value of +2.00V for a chin potentiometer voltage of +10.00V.

2. The position of the chinrest was systematically varied such that the voltage of the chinrest potentiometer was increased from 0.0 to +10.00V in 1.0V increments. For each voltage, the heights of the left- and right-hand sides of the square were measured.

3. Step 2 was repeated a further three times. Fig. 2.6 shows the calibration data. Each point represents the mean of four separate measurements.

Selected readings for chinrest potentiometer voltages between 0.00V and -10.00V were also taken. These confirmed that the degree of trapezoidal distortion was symmetrical about the midpoint of the chinrest's translation path.

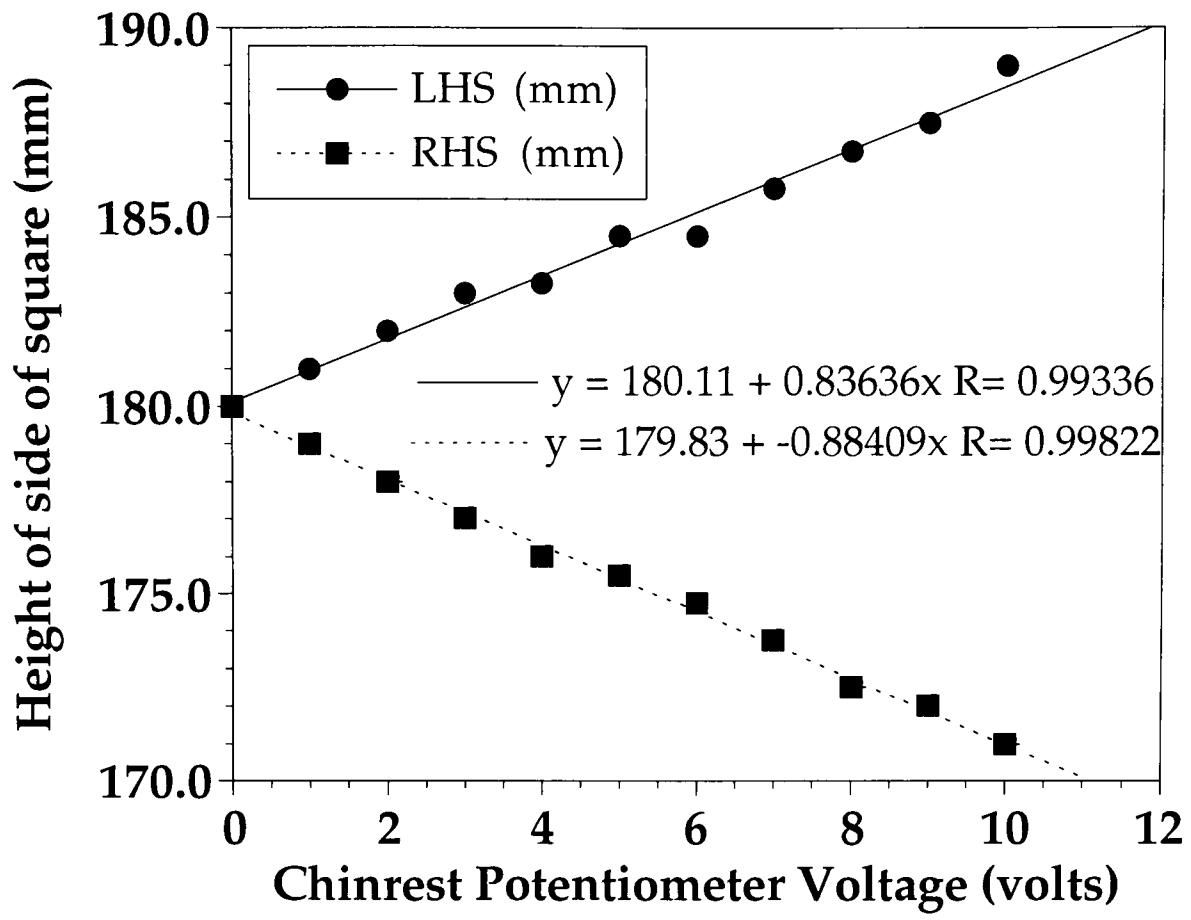


Fig. 2.6: Height of left- and right-hand sides of square as a function of chinrest potentiometer voltage. The amplitude of the W186 was fixed such that it was +2.00V when the voltage of the chinrest potentiometer was +10.00V.

2.1.6 Measurement of magnitude of perceived rotation, using a real paddle

In Chapter 4 of this thesis, the magnitude of perceived rotation of motion parallax surfaces is examined. This was carried out using a matching technique. Observers matched the amount of perceived rotation of the stimulus with that of a real, circular paddle situated at 57 cm. The bottom of the paddle was 1 cm above the top of the oscilloscope screen, and the centres of the paddle and the screen lay on the same vertical axis when the viewing distance (observer to screen) was 57 cm. The paddle was 7 cm in diameter. It was constructed of balsa wood, and its surface was covered with patterned paper. The pattern was a 50% density, irregular textured pattern of 'blobs', each of which subtended 1 - 2°. The paddle rotated about its vertical axis. The rotation of the paddle was linked to the head movement of the subject so that as the head moved to the right, the paddle rotated in a clockwise direction, and as the head moved to the left, the paddle rotated in an anti-clockwise direction. The extent of rotation was controlled with a hand-held potentiometer.

The paddle was controlled by a modified output of one of the pen

motors of an EEG recording unit (Devices M19 Recording Systems). Since this apparatus is usually employed to monitor extremely small changes in electrical conductance, the pen motors which record such changes are extremely sensitive. As such, they are ideal for their modified task since they allowed fine adjustment of the slant angle of the paddle. A circuit diagram illustrating the incorporation of the pen motor output into the general apparatus is shown below in Fig. 2.7.

Experiments 18 and 19 (Chapter 6) examine the perceived rotation of a structure rotating in the same direction as the observer's head movement. The trapezoidal transformations required for this simulation are in the opposite direction to those previously described: instead of a size increase in the vertical edge nearer to the observer, an increase in the size of the further edge is needed. In order to produce this transformation a two-way switch was also incorporated into the circuit.

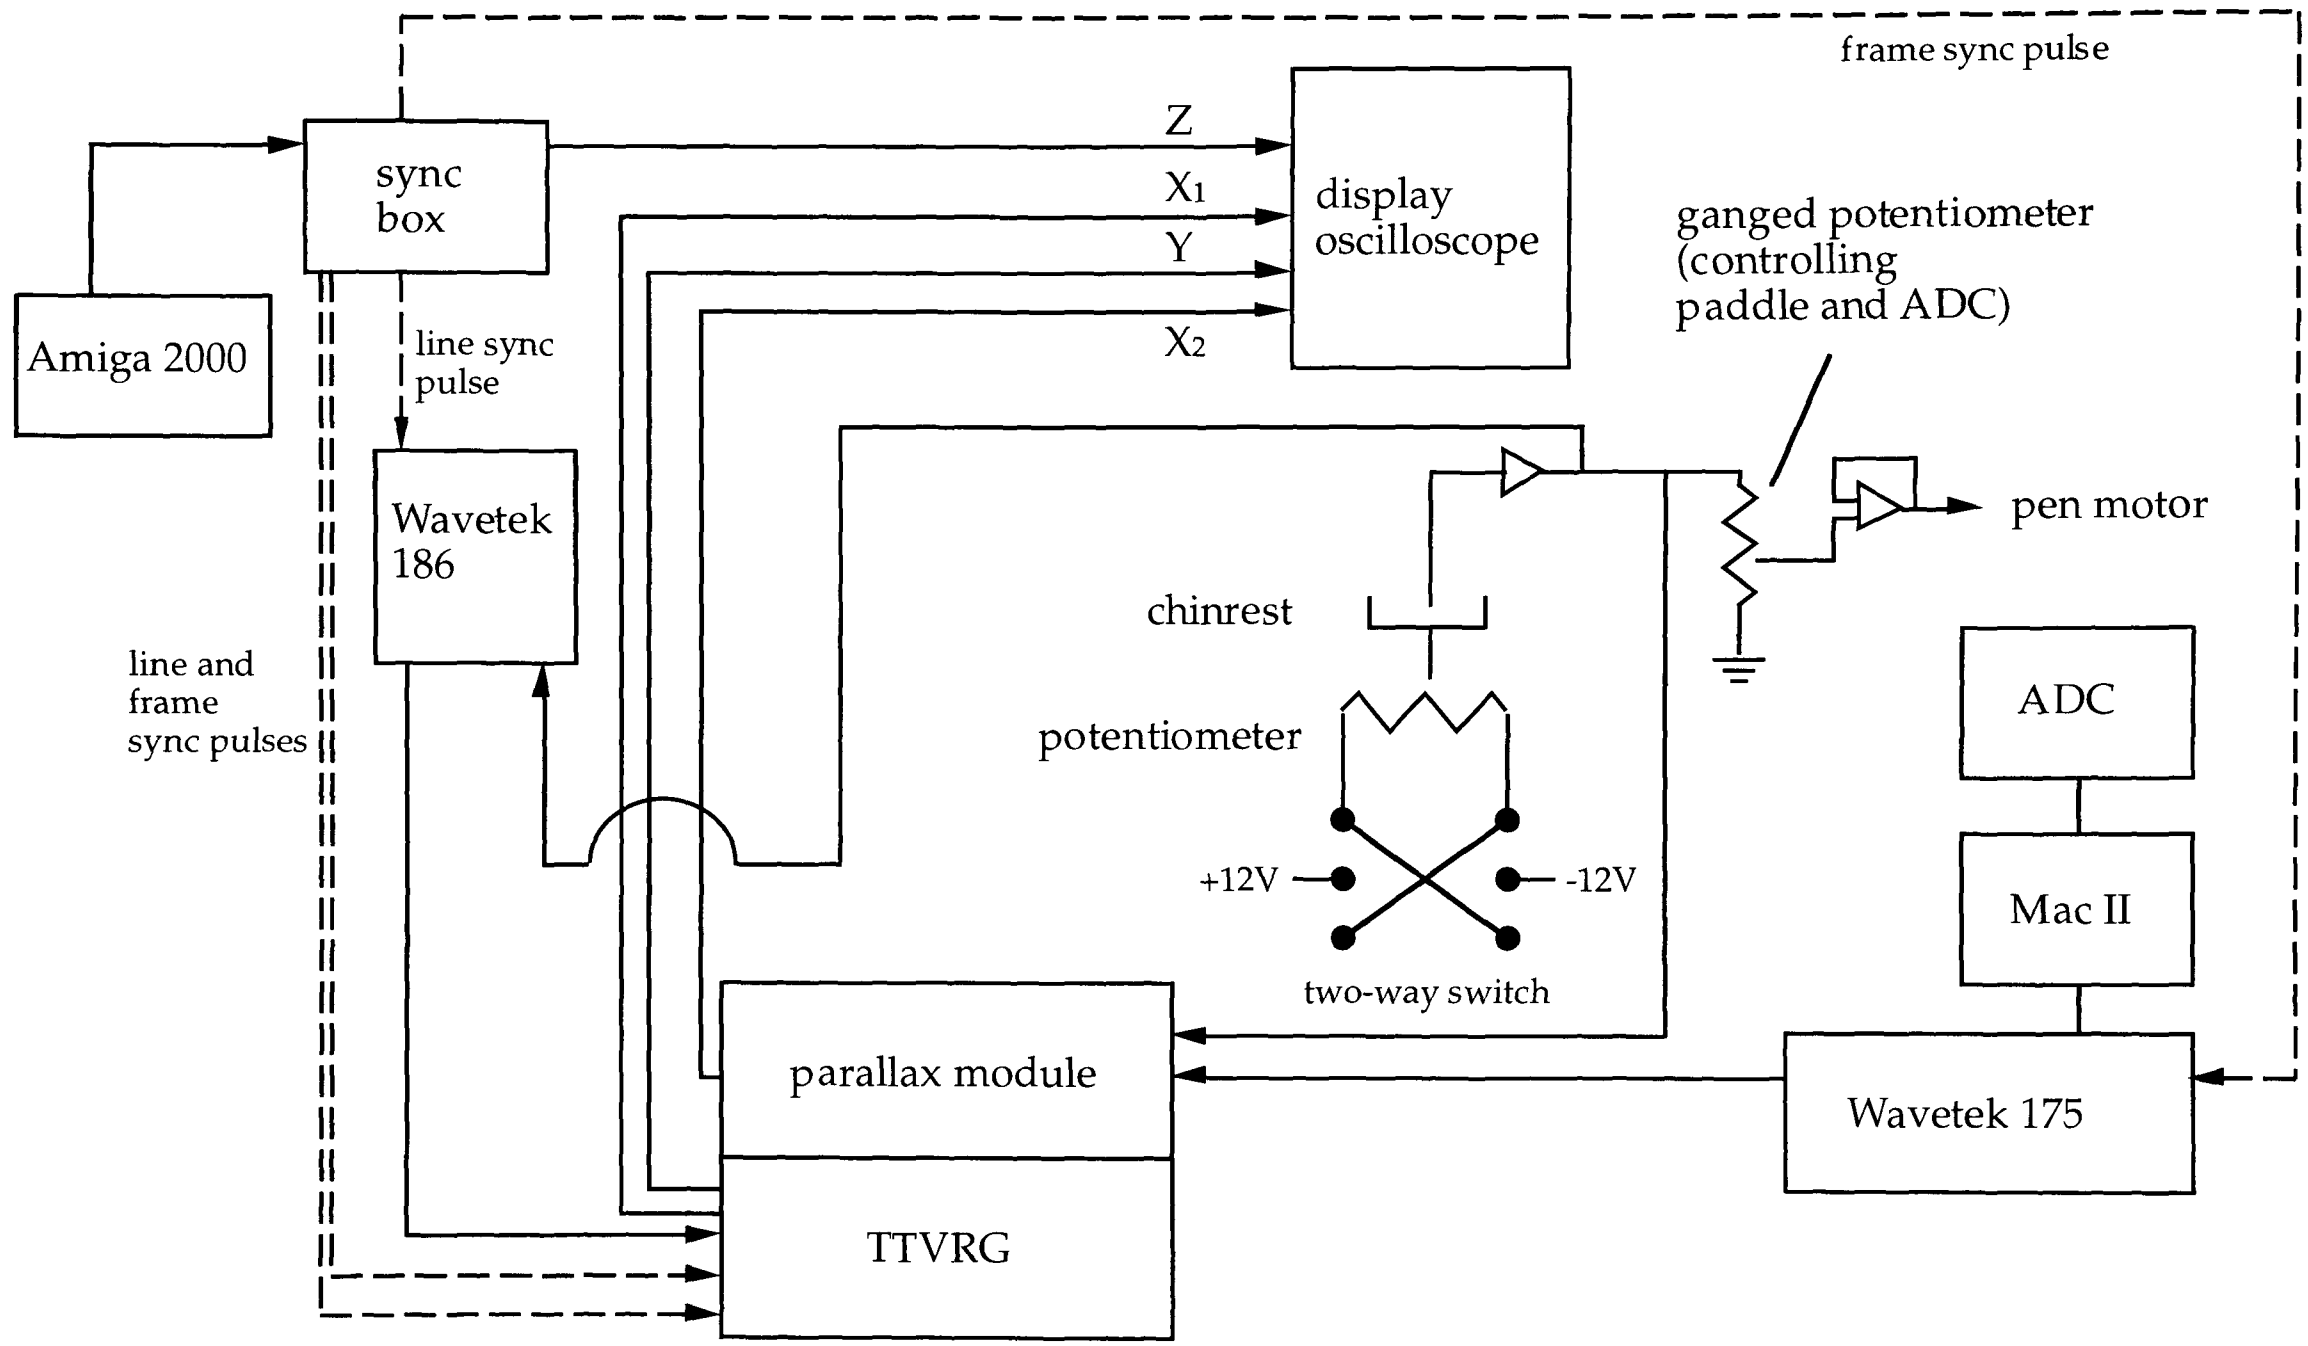


Fig. 2.7: Representation of the modified apparatus.

2.1.7 Calibration of the paddle used for rotation matching

The relationship between the rotation angle and the potentiometer reading was determined using a laser and basic trigonometric principles. A thin piece of glass was temporarily affixed to the paddle surface so that the laser's light could be reflected. For various systematic combinations of chinrest position (measured in terms of the voltage across the chinrest potentiometer) and paddle angle, the variable resistance of the potentiometer controlling rotation angle was measured (Fig 2.8).

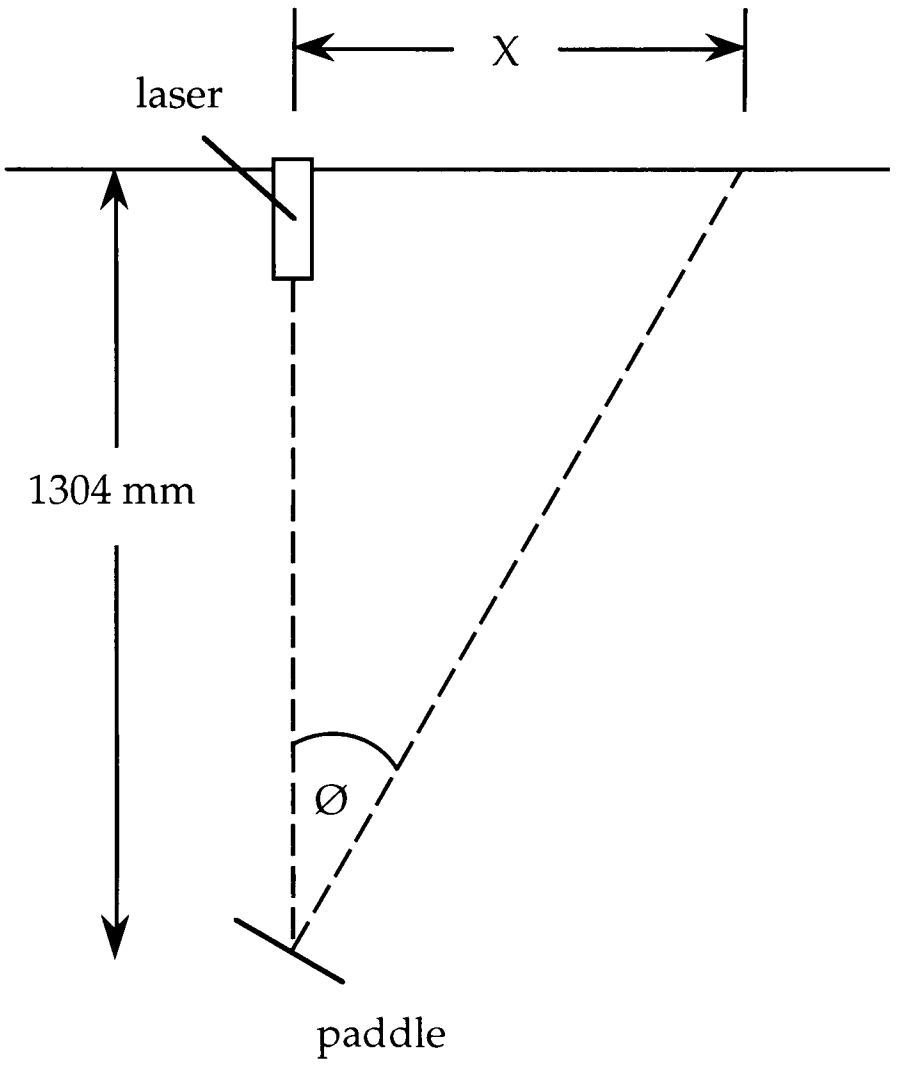


Fig 2.8: Calibration of the paddle angle.

1. The chinrest was fixed in position such that the voltage across the chin potentiometer was fixed at +10.00V.

2. The rotation angle of the paddle was varied such that the distance X increased from 0 to 608.5 mm in increments of 60 mm (the final increment was 8.5 mm).

3. For each value of X, the corresponding ADC output (0-255) was monitored.

4. Steps 2 and 3 were repeated a further three times. The mean of each of these sets of four readings was calculated.

The process was carried out again with the chinrest potentiometer voltage fixed at -10.00V.

The rotation angle of the paddle, Ø, was calculated by simple trigonometry. Fig. 2.9 shows the rotation angle as a function of the ADC reading. Data from the chinrest potentiometer voltages of +10.00V and -10.00V are shown on the same graph, resulting in two rotation angles (positive and negative) for each ADC value.

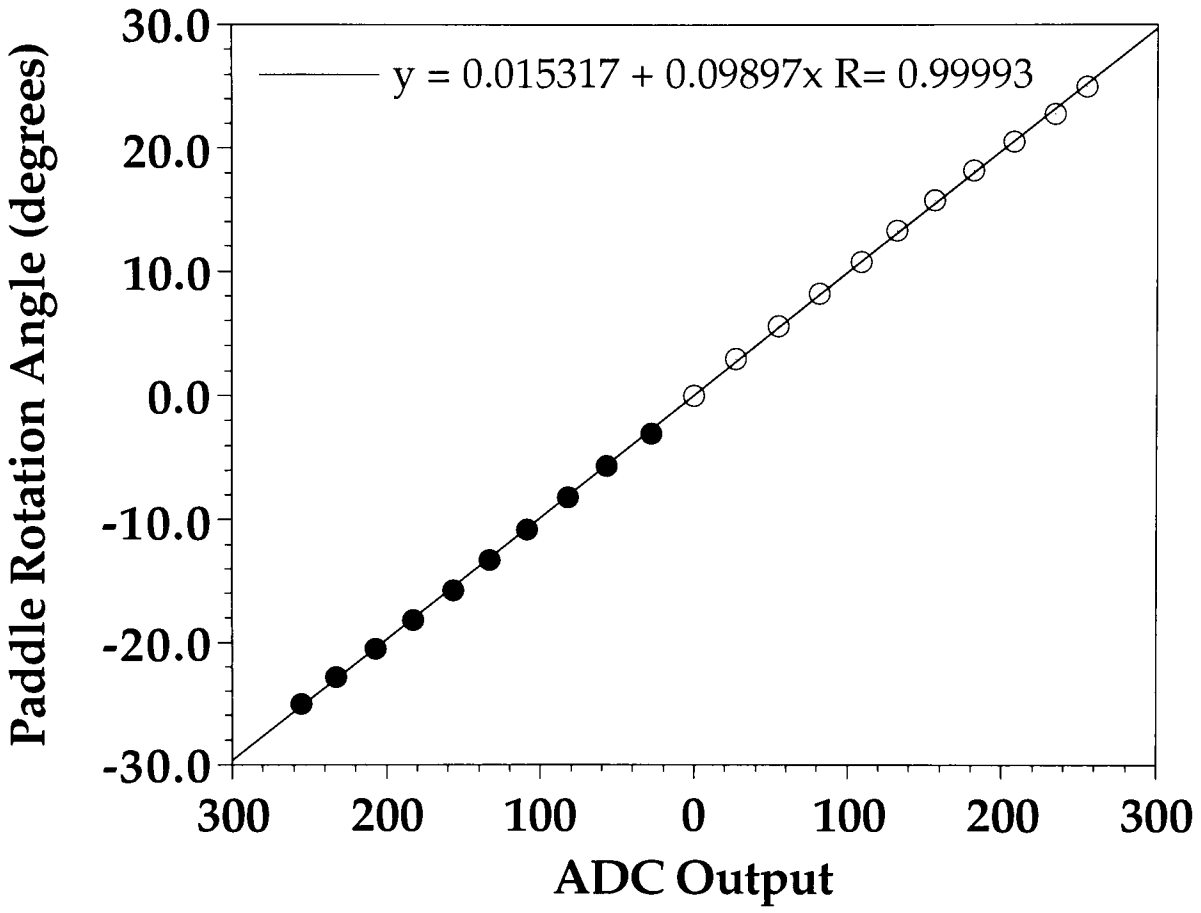


Fig. 2.9: Paddle rotation angle as a function of ADC output. The open and closed symbols represent data for chinrest potentiometer voltages of +10.00V and -10.00V respectively.

In order to measure the linearity of the paddle rotation angle as a function of the chinrest position, a second set of calibration data was collected.

1. The voltage across the potentiometer controlling the paddle rotation angle was set to its maximum value.

2. The voltage across the chinrest potentiometer was varied (by moving the chinrest) so that the voltage increased from 0.00V to +10.00V in increments of 1.00V. For each voltage, distance X was measured, thereby allowing the calculation of the paddle rotation angle.

3. Step 2 was repeated for chinrest potentiometer voltages of 0.00V to -10.00V.

4. Steps 2 and 3 were repeated a further three times. The mean of each of these sets of four readings was calculated.

Fig. 2.10 shows the paddle rotation angle as a function of the chinrest potentiometer voltage.

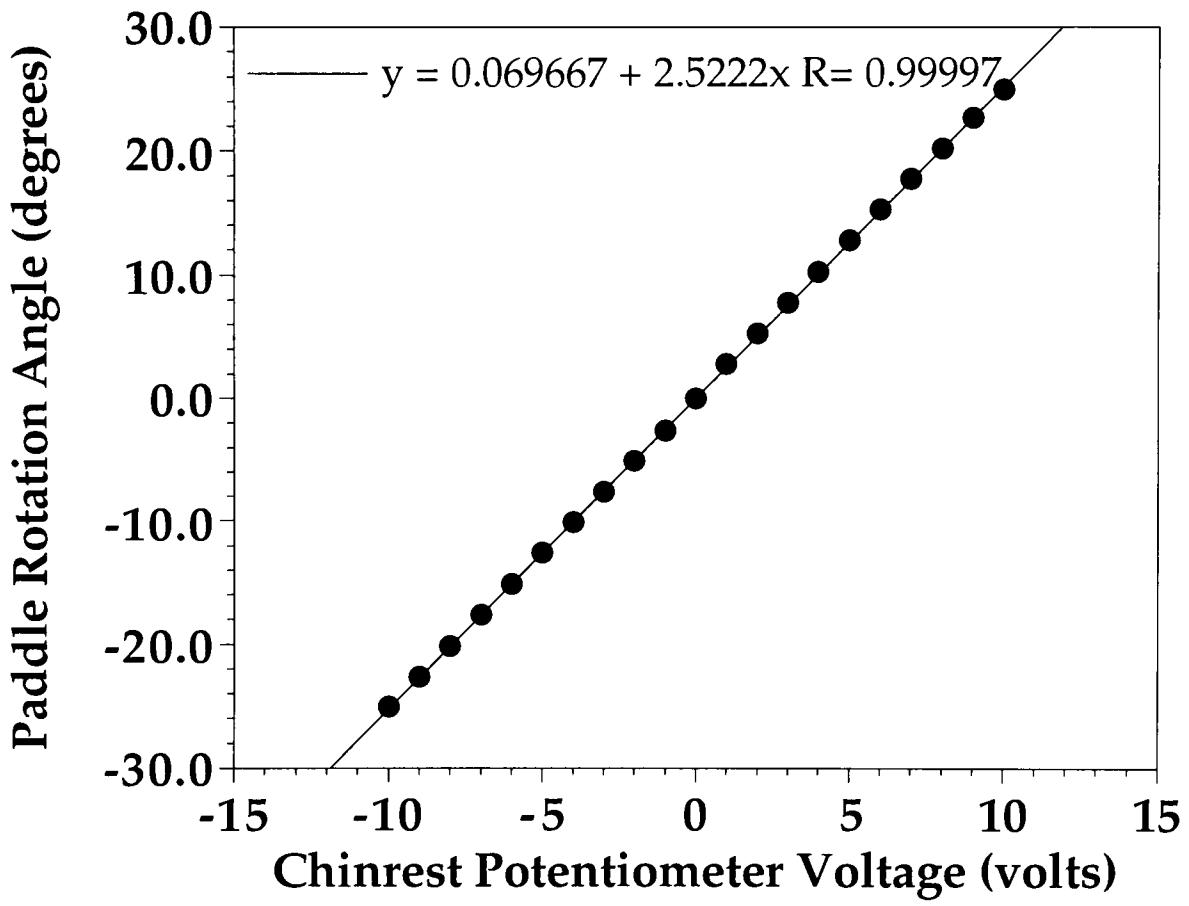


Fig. 2.10: Paddle Rotation Angle as a function of Chinrest Potentiometer Voltage.

2.2 Software

The stimuli were generated on the Amiga 2000 using the Director package. This is a programming package which is specifically designed to generate and present visual displays. A second package, Grabbit, was used to save the screen images in their pictorial format. The programs used for

controlling the experimental sequence were written using Microsoft QuickBASIC.

2.3 Subjects

The subjects who served in this thesis were aged 24 - 27 and were experienced psychophysical observers. All except the author were naïve to the purposes of the experiments. All had normal, or corrected to normal, eyesight.

2.4 Psychophysical Procedures

The psychophysical procedures used in this thesis were the Method of Constant Stimuli and the Method of Adjustment.

2.4.1 Method of Constant Stimuli

A series of stimuli at various levels are presented to the observer, repeatedly and in a random order. The observer makes a simple response to each stimulus. The proportion of responses of a given type is plotted as a function of stimulus level, and a cumulative normal function is fitted to the data using Probit analysis (Finney, 1971). The curve is termed the psychometric function. The function is used to determine the point at which the observer is performing at a specific level. This point may be taken as a sensory threshold for instance, which is generally arbitrarily defined as the 75% correct level. Alternatively it may correspond to a Point of Subjective Equality (PSE), in which case it is the 50% point.

2.4.2 Method of Adjustment

In this psychophysical procedure the observer adjusts the signal strength until the stimulus meets specific perceptual criteria which have been set by the observer. Instances in which this method might be used include the determination of the point at which depth is just detectable in a motion parallax display (e.g. Rogers and Graham, 1982), or the adjustment of a match stimulus so that a given attribute is of the same

magnitude as that of another stimulus, typically known as the test (e.g. van Veen and Werkhoven, 1996; Chapters 4 - 6 of this thesis). The main feature of this method is that the stimulus changes are under the control of the observer which, under some circumstances, means that there is difficulty in maintaining constant conditions during threshold measurement (Gescheider, 1997).

In ascending methods of adjustment the signal strength is slowly increased until the relevant perceptual criteria are met. In descending methods of adjustment the signal strength is slowly decreased until the criteria are met. The former can lead to over-estimations of the point at which this occurs, while the latter may lead to under-estimations. For this reason a preferred method is to make initial coarse adjustments so that the point is bracketed. Following this, increasingly fine adjustments are made.

2.4.3 Note

Note that neither the Method of Constant Stimuli nor the Method of Adjustment leads to bias-free measurements. In some experiments of this thesis these psychophysical methods were employed to determine the PSE between two qualitatively different percepts. For instance, in Chapter 3, subjects were asked to decide whether a motion parallax-defined surface appeared to be either stationary or rotating (see Section 3.4). An objective measure of the PSE in this case is impossible, because observers' judgements are necessarily subjective. Subjects were, however, instructed to attempt to maintain consistency in their judgements, keeping their own perceptual criteria for a given response constant.

Chapter 3:

The Transition Point between a Rigid, Stationary Surface and a Rigid, Rotating Surface; Perceptual Zones*

3.1 The concept of perceptual zones divided by thresholds

The 2-D relative motion underlying a motion parallax stimulus can give rise to a number of different percepts, depending on parameters such as waveform type, viewing distance, amplitude of relative motion etc. For example, when the amount of relative motion is sub-threshold, subjects will obviously not perceive any depth or motion, and the dots comprising the display will appear to be stationary and 2-D, lying in the plane of the monitor. At higher amplitudes subjects typically perceive a stationary 3-D surface, the shape of which is dependent upon the type of waveform of the relative motion (e.g. Rogers and Graham, 1979; Ono and Steinbach, 1990). Depth thresholds for sinusoidally corrugated surfaces have been found to lie between 100 arc sec equivalent disparity (0.05 cpd) and 40 arc sec equivalent disparity (1.6 cpd) (Rogers and Graham, 1982).

As amplitude is increased beyond threshold, subjects generally perceive increasing amounts of depth (e.g. Rogers and Graham, 1979; Ono, Rivest and Ono, 1986; Ono and Steinbach, 1990). Studies have additionally reported the perception of motion of the rigid 3-D surface. This generally takes the form of either a 'rocking' motion — the structure appears to rotate about a vertical axis — (e.g. Ono et al., 1986; Ono and Steinbach, 1990) or a shearing motion — different parts of the structure, whilst remaining locally rigid, appear to move in opposite directions — (Ono and Ujike, 1993; Ichikawa and Ono, 1996). The point at which a parallax surface

---

* Some of the work of this chapter was presented at ARVO, 1996

ceases to appear stationary and starts to rotate or shear has been termed the 'motion threshold' (Steinbach, Ono and Wolf, 1991; Ono and Ujike, 1993; Ichikawa and Ono, 1996). Studies which have examined this threshold are discussed further in Section 3.2.

The various percepts which may arise from a parallax stimulus have been considered in terms of perceptual 'zones', which are divided by perceptual thresholds (e.g. Saida and Ono, 1989; Ono and Ujike, 1993). For instance, the depth threshold divides the zone in which no depth is perceived from that in which a rigid, stationary 3-D structure is perceived. The 'motion threshold' divides the zone within which a rigid, stationary 3-D structure is perceived from that within which a rigid, moving (rotating or shearing) 3-D structure is perceived. Experiments 1 - 4 examine the factors upon which this threshold depends. Experiment 5 discusses perceptual zones in more detail, and determines the thresholds dividing five such zones.

3.2 The 'motion threshold'

Evidence suggests that a variety of factors influence the perceived motion of parallax structures.

3.2.1 The effect of waveform type

Ono and Steinbach (1990) presented observers with parallax surfaces containing motion gradients with sine, triangle, sawtooth or square waveforms. It was found that the greatest amount of motion was perceived in the square wave surface, followed by the ramp, triangle and sine waveforms respectively. However, it was also found that the nature of the motion was dependent on the waveform. The square wave appeared to undergo 'shearing' motion, whereas the other waveforms appeared to rotate about a vertical axis.

3.2.2 The effect of relative motion amplitude

As noted above, studies have reported that as the amplitude of relative motion increases, a greater number of observers perceive the parallax

structure to rotate or to shear.

In the study of Ono and Steinbach (1990) (see Section 3.2.1), the parallax surfaces contained relative motion amplitudes of 12 or 113 arc min equivalent disparity. Observers made depth judgements and indicated whether or not the dots on the surface appeared to translate, and if so, by how much. It was found that the extents of both perceived motion and perceived depth were greater for the higher relative motion amplitude.

Braunstein and Andersen (1981) presented stationary observers with translating, horizontally-oriented dihedral angles formed by two slanted planes meeting at a horizontal line (see Fig. 1.5). The angle formed by the intersecting planes was either the nearest or the furthest part of the display. The maximum dot speed was 2.6, 5.2 or 10.4 deg sec^{-1}. The minimum dot speed was either half the maximum speed or zero. Subjects were asked to indicate the perceived depth order and to classify the amount of perceived rotation in the display as "no rotation", "some rotation", or "mostly rotation". The mean number of observers perceiving "mostly rotation" increased with increasing maximum dot speed, and was greater when the minimum speed was zero than when it was non-zero.

3.2.3 The effect of viewing distance

Ono et al. (1986) investigated depth perception as a function of motion parallax and absolute distance information. Observers were asked to indicate the amount of perceived depth in sinusoidal parallax surfaces at viewing distances of either 40, 80, 160 or 320 cm, and to indicate whether or not the surface appeared to rotate. Three levels of distal parallax were used. It was found that as viewing distance increased, the number of observers perceiving the parallax surface to rotate increased also. At each viewing distance, as the amount of parallax increased the number of observers perceiving rotation also increased.

3.2.4 The effect of head movement velocity

Ono and Ujike (1993) and Ichikawa and Ono (1996) investigated the 'motion' threshold as a function of head movement velocity. The former study investigated velocities between 0.125 and 4.0 cm sec^{-1}, and found that

thresholds dropped with a slope of -1 for head velocities up to 0.5 cm sec^{-1} but remained constant for velocities greater than this. The latter study examined head movement velocities between 0.5 and 16 cm sec^{-1} and found that motion thresholds fell with a slope of approximately -1 throughout this entire range. Ujike and Ono (1992) investigated depth thresholds as a function of the amplitude (5 - 30 cm) and frequency (0.083 - 1.25 Hz) of head movement and found that for the higher head frequencies, thresholds did not differ as a function of head movement amplitude but for lower head frequencies thresholds were higher for the small amplitudes than for the large. The above studies used surfaces with square wave depth modulations.

3.2.5 The effect of type of head movement

Steinbach, Ono and Wolf (1991) examined depth and motion thresholds as a function of the direction in which the head moved (horizontal or vertical) and the type of head movement made (translation or rotation). The rotational movements were about either a horizontal (a nodding, 'yes' movement) or a vertical (a shaking, 'no' movement) axis. The stimulus was a sinusoidally corrugated surface. A descending method of adjustment was used to obtain the thresholds. In separate conditions, subjects decreased the amount of relative motion in the display until depth was no longer perceived (depth threshold) or until motion was no longer perceived ('motion' threshold). It was found that the direction, but not the type, of head movements affected the depth thresholds. Thresholds were lower when the head motion was horizontal. The motion thresholds were affected by neither the type nor the direction of head motion.

3.2.6 The effect of observer movement

There is some evidence to suggest that thresholds for perceiving rotation may be partly dependent on whether the observer is stationary or moving. In the study by Ono and Steinbach (1990), observers viewed motion parallax displays in which the relative dot motion was either head-linked or externally generated. Two extents of relative motion were used: 12 and 113 arc min equivalent disparity. The relative dot motions had gradients of sine, square, triangle or ramp waveforms. It was found

that, for all waveforms, less rotation was perceived in the head movement condition than in the no head movement condition.

Ichikawa and Ono (1996) investigated motion thresholds under conditions of observer movement and no head movement, for head/stimulus velocities between 0.5 and 16 cm sec^{-1}. The motion threshold was slightly lower in the latter condition throughout almost the entire range of velocities.

Rogers and Graham (1979) found that the magnitude of perceived depth was consistently lower when motion parallax was created by the translation of the oscilloscope while the observer remained stationary than when the parallax was linked to the observer's head motion (Fig. 3.1).

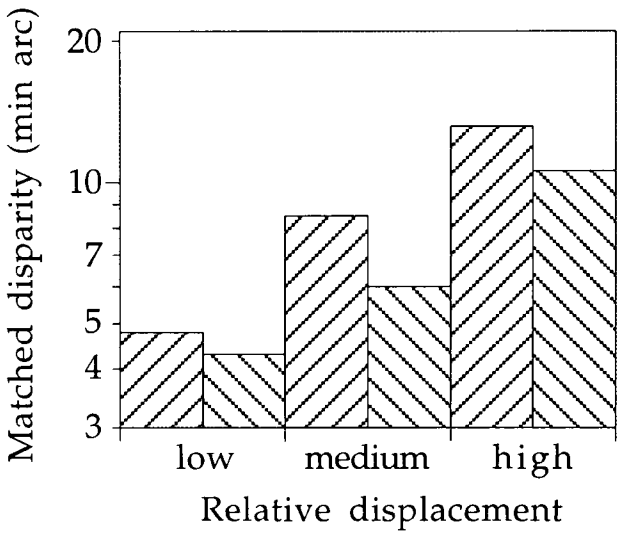


Fig. 3.1: Matched stereoscopic depth as a function of relative motion in motion parallax displays. The left-hand bars represent results obtained with self-produced movement and the right-hand bar those obtained under externally produced parallax conditions. Taken from Rogers and Graham (1979, Fig. 5, p. 132).

The perception of rotation was not investigated in this study. Other studies have reported that a decreased amount of depth is often perceived in conjunction with perceived rotation (Ono, et al., 1986; Rogers and Collett, 1989; Ono and Steinbach, 1990; Sakuri and Ono, 1996). It may be the case that had rotation been examined, its perceived magnitude would have been greater in the condition in which the parallax was externally produced than in the condition in which it was linked to the head motion of the observer.

Ono and Steinbach (1990) suggested that the increased propensity to perceive rotation when the observer is stationary than when they are

moving is due to the operation of 'location constancy' (Graham, 1951; Ittelson, 1960; Mack, 1986). When an observer moves in the real world, the retinal motion resulting from the observer's own movement must be compensated for if accurate information about the surrounding environment is to be gained. For instance, turning or nodding head movements cause displacements of the environment relative to the eye which are identical to those which would be produced by small rotations of the environment about the viewer. However, if the environment moved it would be perceived to do so, whereas when the observer turns their head the environment is perceived to remain stationary. This is because proprioceptive information specifies that in the latter case the head is moving and in the former case it is not.

Wallach and colleagues have carried out extensive research on the perception of a stable environment in the face of observer movement (e.g. Wallach and Flaherty, 1975; Wallach, Bacon and Schulman, 1978; Whipple and Wallach, 1978; Wallach, 1985a, 1985b). One instance of compensation, namely that for the relative rotations of objects in the environment when we move forward, will be considered here since it has specific relevance to motion parallax paradigms.

When we move forwards, a stationary object on our left will rotate in a clockwise direction, relative to the eye. Our compensatory mechanisms allow us to discount this rotation however, and the object is perceived as stationary. Wallach (1985) investigated the accuracy of the compensation process by examining observers' percepts when the object did not remain stationary as they passed it, but rather rotated very slightly. A variable ratio transmission was attached to the ceiling, and to its output shaft was attached a small patterned sphere. When the observer moved back and forth past the object, relative rotation of the object occurred (see Fig. 3.2). The observer's position relative to the object was transmitted to the variable ratio transmission, "... which in turn could make the test object rotate in either direction and in any proportion of the relative rotation caused by the observer's changing position. It could thereby cause that relative rotation to increase or decrease" (Wallach, 1987, p. 6). It was found that observers were not particularly good at detecting whether or not the ball was objectively rotating: it could rotate through an angle of 40% of its relative rotation before observers perceived that it was rotating on its own.

In effect, this means that under these circumstances, observers are not very adept at disentangling subjective rotation and objective rotation: they can not distinguish between rotation that arises purely as a result of their own changing angle with respect to the object, and rotation that arises because the object in question is actually rotating about one of its own axes.

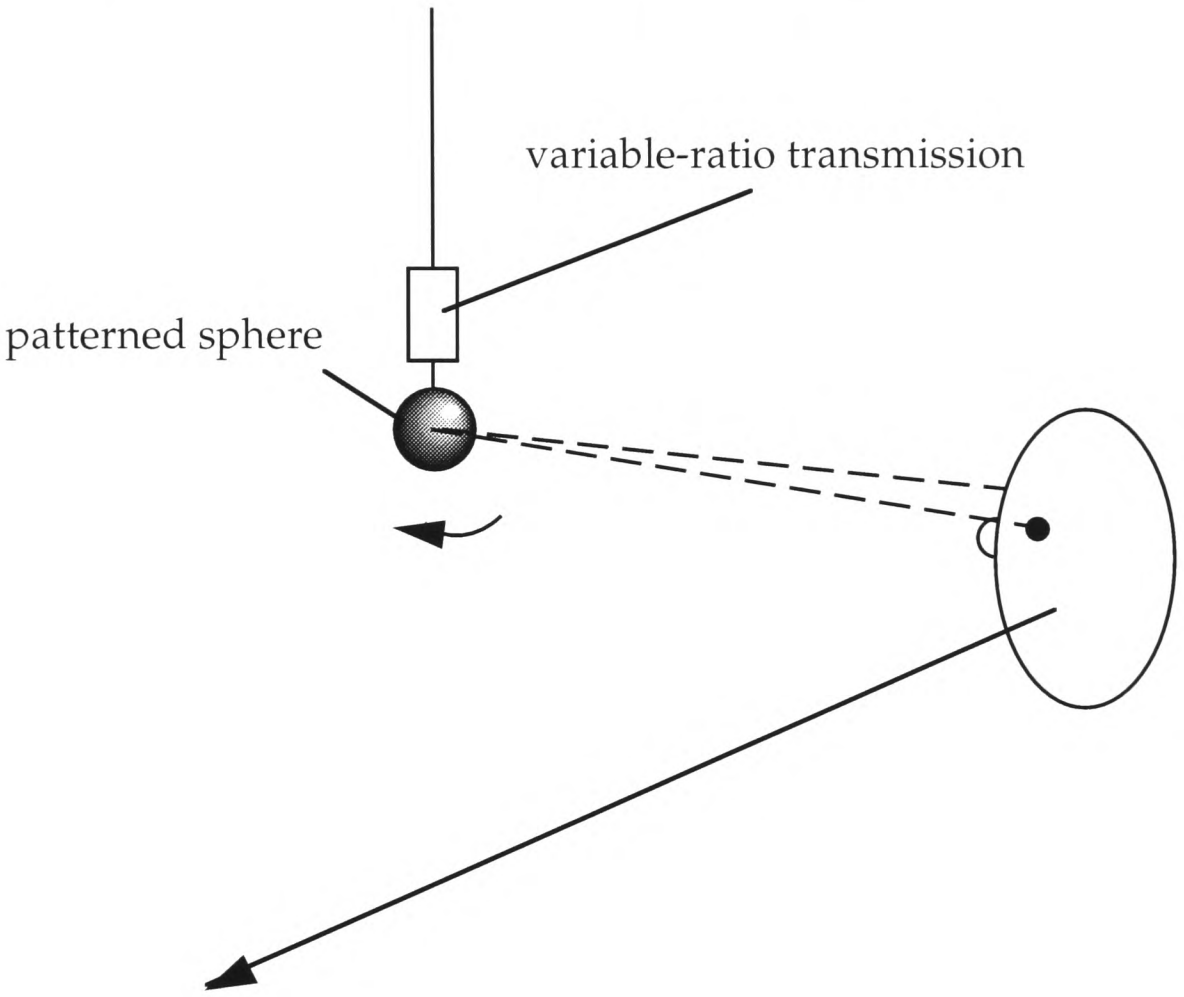


Fig. 3.2: The observer's transverse movement was transmitted to the input shaft of a variable-ratio transmission, which turned a patterned globe. The globe could rotate in either direction, and the extent of rotation was determined by the setting of the transmission. The apparatus was designed to investigate the extent of departure required from the normal rotation of an object with respect to the passing observer before the observer is aware of the object's rotation. Taken from Wallach (1985, p. 95).

It is as though the moving observer assumes more readily than he should that the world is stationary: when an object is stationary it is perceived as such, but when the object is actually rotating, it is still often perceived as stationary. Wallach terms the range of relative environmental displacements which result in the perception of the environment as stationary the 'immobility range' (e.g. Wallach, 1985; 1987).

The immobility range is a real world example of the effects of the ambiguity of retinal motion: the motion could have arisen *either* from the

relative rotation of the object due to the viewer's own movement (due to the fact that different parts of the objects are at different distances from the observer and thus produce relative motion — a single point of light, for instance, could not appear to rotate) *or* from the objective rotation of the object, or from a combination of these sources.

Wallach's experimental paradigm is analogous to the situation in which a horizontally translating observer views an object before him, and has thus been applied to situations involving motion parallax stimuli presented on a frontal surface (Ono et al., 1986; Ono and Steinbach, 1990; Steinbach et al., 1991).

Observers' relative insensitivity to objective rotation when it is combined with self-motion suggests the existence of a bias against the interpretation of retinal motion as arising from object movement, when observer movement is occurring.

This bias found with real objects may underlie the tendency of observers to interpret the relative motion of head-linked motion parallax displays as depth rather than rotation. When motion parallax displays are viewed by *stationary* observers the immobility range obviously does not come into play, since there is no observer-generated retinal motion to be compensated for. The bias against the perception of rotation observed under conditions of observer movement is not in operation and the relative motion is 'free' to be interpreted as rotation. Ono and Steinbach (1990) proposed that '... the afferent signals from the retinal motion in the head-movement and no-head-movement conditions are coded by the same motion detectors, but at some stage the visual system must convert the motion signal to create a perception of a stationary surface when the retinal motion is produced by head movement'.

3.3 Experiments of this chapter

The above studies suggest that the point at which a parallax structure appears to start to rotate or to shear is dependent upon a variety of factors. These include viewing distance, head movement velocity, amplitude of relative motion and whether or not the observer is moving. However,

some difficulties exist with these conclusions.

Firstly, the studies of Ono et al. (1986), Ono and Steinbach (1990) and Braunstein and Andersen (1981) were not concerned with the 'motion threshold' itself. Thus the results are of a rather qualitative nature, and firm conclusions regarding the precise relationship between the threshold and the manipulated variables cannot be reached.

Furthermore, the suggestion of Ono et al. (1986) that observers are increasingly likely to perceive rotation as viewing distance is increased is not necessarily true. The distal stimulus of that experiment remained constant as viewing distance increased. Corrugation frequency and stimulus size were therefore respectively doubled and halved with each doubling of viewing distance. Consequently, the observed increase in the occurrence of perceived rotation could be at least partly attributable to variations in either of these two variables.

Ono and Steinbach's (1990) finding that square, ramp, triangle and sine waveforms could be ordered with respect to their relative amounts of perceived motion must be tempered by the observation that the type of motion was not identical across waveforms. The type of motion to be judged must be specified at the outset of the experiment if legitimate comparisons between the data of different waveforms are to be made.

Some of the studies investigating perceptual thresholds have used unsuitable stimuli. For instance, each band of the square wave surfaces of Ichikawa and Ono (1996), Ono and Ujike (1993) and Ujike and Ono (1992) consisted of vertically-oriented sinusoidal luminance modulations. A potential problem is that the relative positions of these bands could have been determined purely on the basis of position cues. It is preferable to use random dot patterns, which do not provide such cues.

The first five experiments of this chapter aim to uncover quantitative relationships between the 'motion threshold' and various factors. Experiments 1a and 1b examine the effect of corrugation frequency upon the threshold. Experiment 2 examines the effect of waveform. Experiment 3 examines the effect of stimulus size. Experiment 4 examines the effect of viewing distance for proximally equal stimuli, thus eliminating the

possible confounding variables of Ono et al.'s (1986) study.

The term 'Transition Point'

The nature of the motion perceived within the perceptual zone above the 'motion threshold' is dependent on the waveform. As noted above, sinusoidal structures appear to rotate (Ono and Steinbach, 1990; Ono et al., 1986) while square wave structures appear to undergo 'shearing' motion (Ono and Steinbach, 1990; Ono and Ujike, 1993; Ichikawa and Ono, 1996). In this thesis the stimuli are generally sinusoidal. The point at which rotation is perceived will henceforth be termed the 'Transition Point'.

3.4 Methods

The following procedure was carried out for each of the following experiments, unless otherwise stated. On each trial, observers decided whether the viewed surface appeared to be rigid and stationary, or rigid and rotating. The method of constant stimuli was used to determine the point at which a 'rigid, rotating' response was equally as likely as a 'rigid, stationary' response. On each trial the relative amplitude of the dot motion was chosen randomly from seven possible values. This set of values had been chosen to ensure that the lowest value was always perceived as a rigid, stationary surface, and that the highest value was always perceived as a rigid, rotating surface.

It will be noted that this method does not lead to an objective measure of the Transition Point. Such a measure is impossible, since observers' judgements are necessarily subjective — the Transition Point is itself a subjective estimate. Subjects were, however, instructed to attempt to maintain consistency in their judgements, keeping their own perceptual criteria for a given response constant (see also Section 2.4.3).

Trials were blocked in groups of 70, corresponding to ten repetitions of each relative motion amplitude. Subjects completed four sessions of ten trials per point per session for all conditions, corresponding to a total of 40 trials at each relative motion amplitude within each condition.

3.5 Experiment 1a: Transition Point as a function of Corrugation Frequency

3.5.1 Introduction

If the Transition Point is dependent upon the maximum motion gradient (MMG) of the surface, then it is predicted that for each doubling of corrugation frequency there will be a halving of the Transition Point. A log-log plot of Transition Point against corrugation frequency would therefore have a gradient of -1. If the Transition Point is dependent only upon the absolute amount of relative motion then it is predicted that the function will have a gradient of zero. Fig. 3.3 illustrates these predictions.

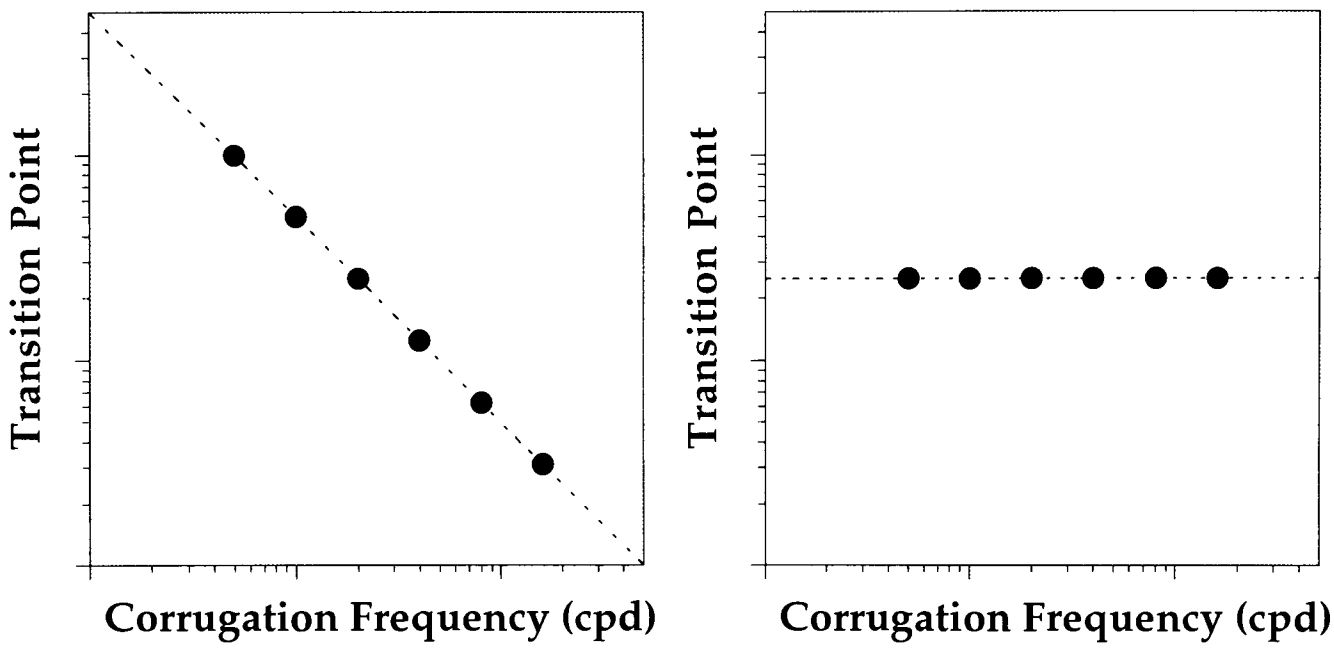


Fig. 3.3: The left-hand graph shows the predicted relationship between Corrugation Frequency and Transition Point if Transition Points are determined by the MMG of the surface. The right-hand graph shows the predicted relationship between Corrugation Frequency and Transition Point if Transition Points are determined by the absolute Peak-to-Trough amplitude of the surface.

3.5.2 Stimuli

The stimulus was a circular 50% random dot pattern of size 20°. It was sinusoidally modulated in depth with one of the following corrugation frequencies: 0.05, 0.1, 0.2, 0.4 and 1.6 cpd.

3.5.3 Results

For each corrugation frequency, the proportion of 'rigid, rotating' responses was plotted as a function of relative motion amplitude. The

best-fitting cumulative Gaussian curve was determined using the Probit technique (Finney, 1971). In all cases the curve fit the data well. The 50% point was taken as the 'Transition Point', the point at which a response of 'rigid, rotating' was equally as likely as a 'rigid, stationary' response. The Standard Errors of the Means (SEMs) of these data and those of subsequent experiments of this chapter were generally very small. Consequently, error bars are not always visible. Where the error bars are visible, they represent ± 1 SEM. Fig. 3.4 shows an illustrative psychometric function.

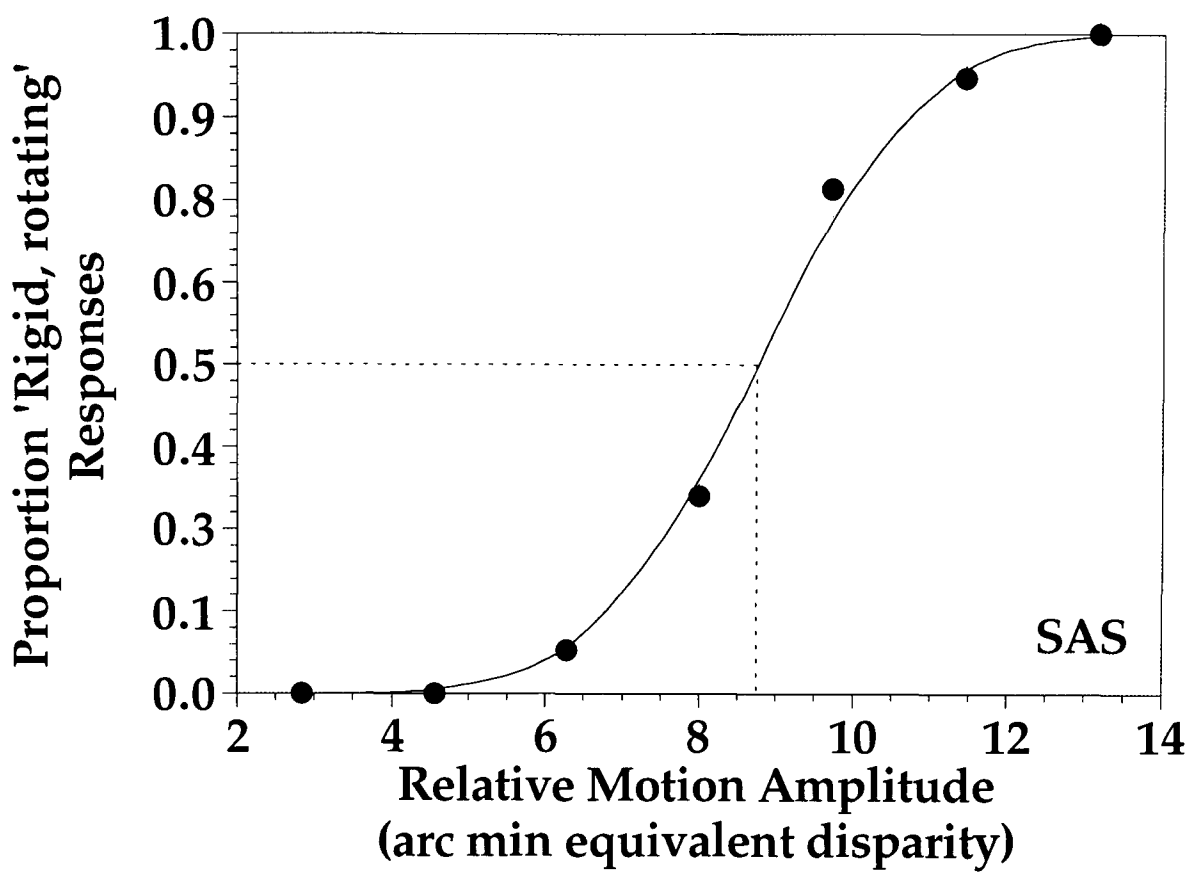


Fig. 3.4: Psychometric function for a 0.05 cpd surface for Subject SAS. The filled circles show the proportion of 'Rigid, rotating' responses as function of the relative motion amplitude. The dotted line shows the 50% point of the best-fitting cumulative Gaussian curve.

Figs 3.5 - 3.7 below show Transition Points (expressed in terms of equivalent disparity; see Section 1.3.4) as a function of corrugation frequency for the three subjects. Over a range of corrugation frequencies the Transition Point fell linearly with increasing corrugation frequency, with a slope of approximately -1. At the lowest frequency, 0.05 cpd, Transition Points were in the region of 10 arc min equivalent disparity for all subjects. For subjects SAS and SJP, increases in corrugation frequency beyond 0.8 cpd did not lead to a further decrease in the Transition Point amplitude, which asymptoted at approximately 1 arc min equivalent disparity. For subject AMJ, the Transition Point function asymptoted at 0.4 cpd and about 2.5 arc min equivalent disparity.

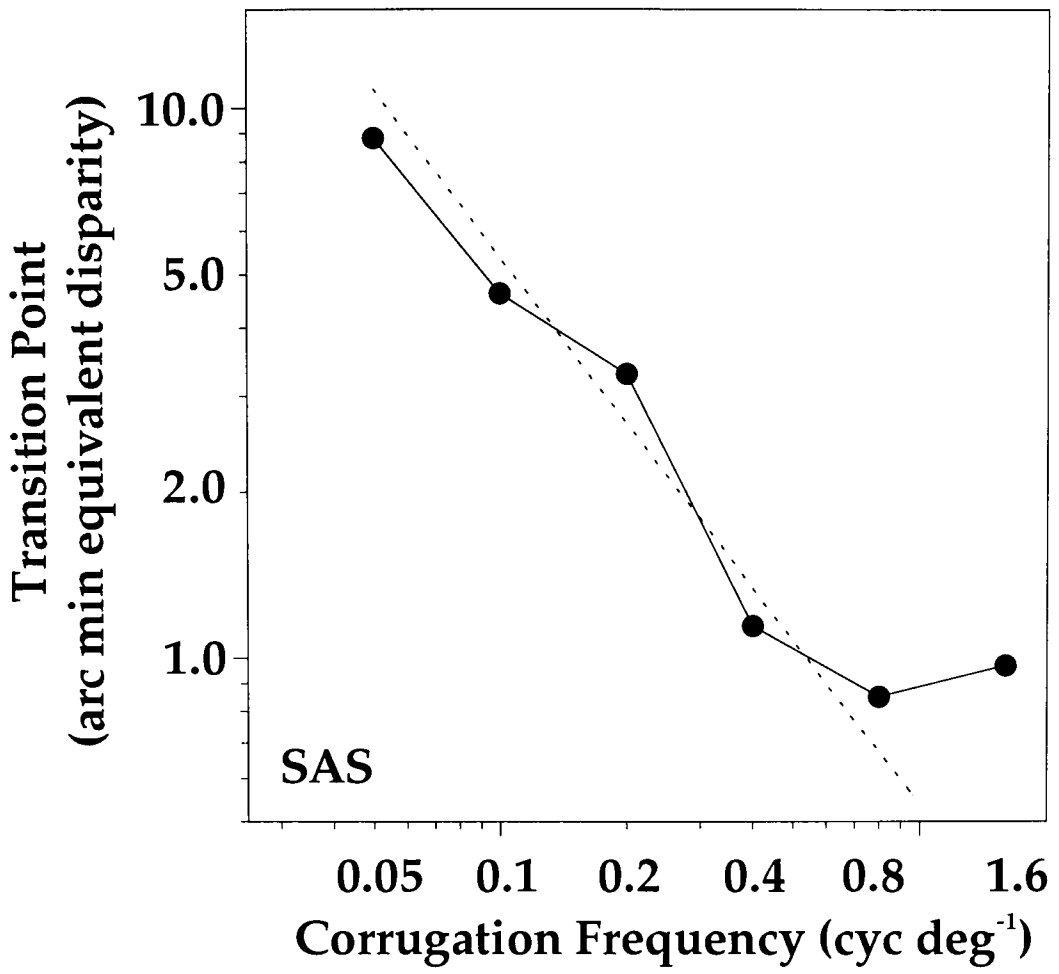


Fig. 3.5: Transition Point as a function of Corrugation Frequency for Subject SAS. The dotted line has a slope of -1.

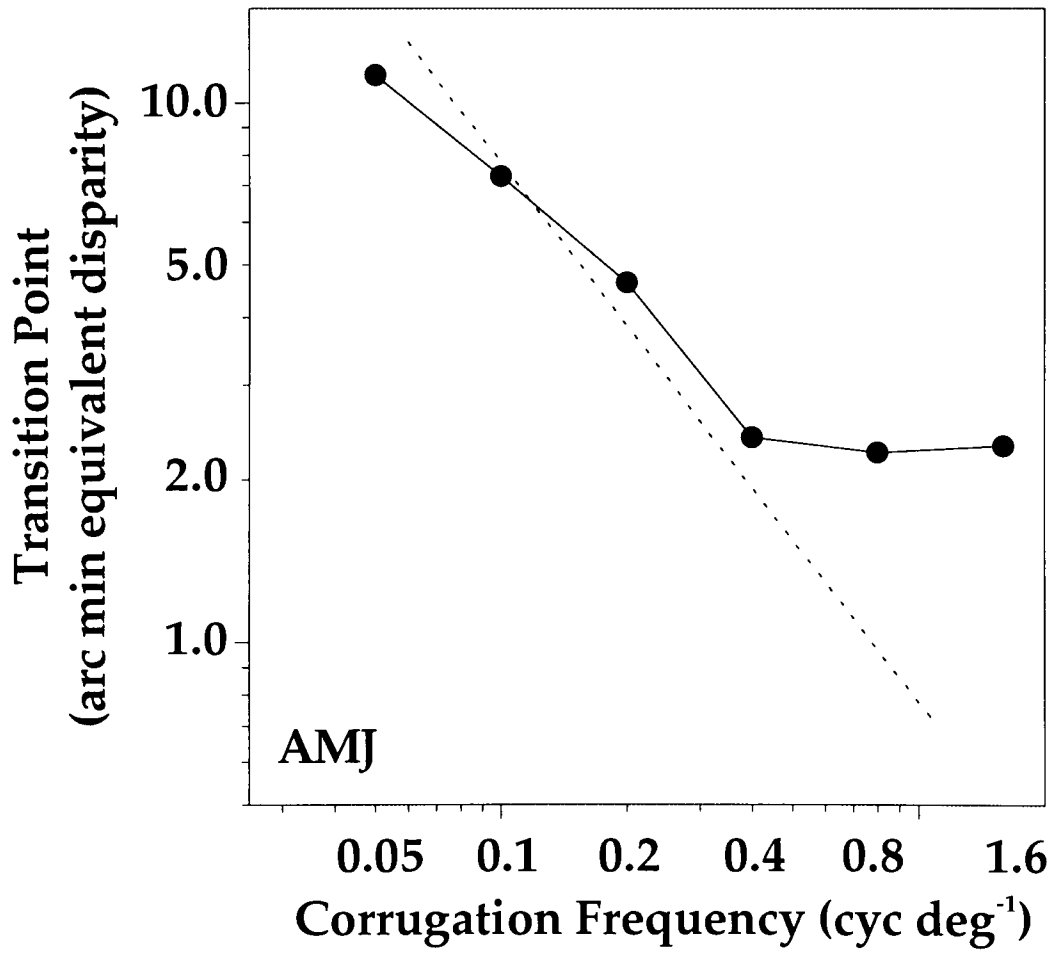


Fig. 3.6: Transition Point as a function of Corrugation Frequency for Subject AMJ. The dotted line has a slope of -1.

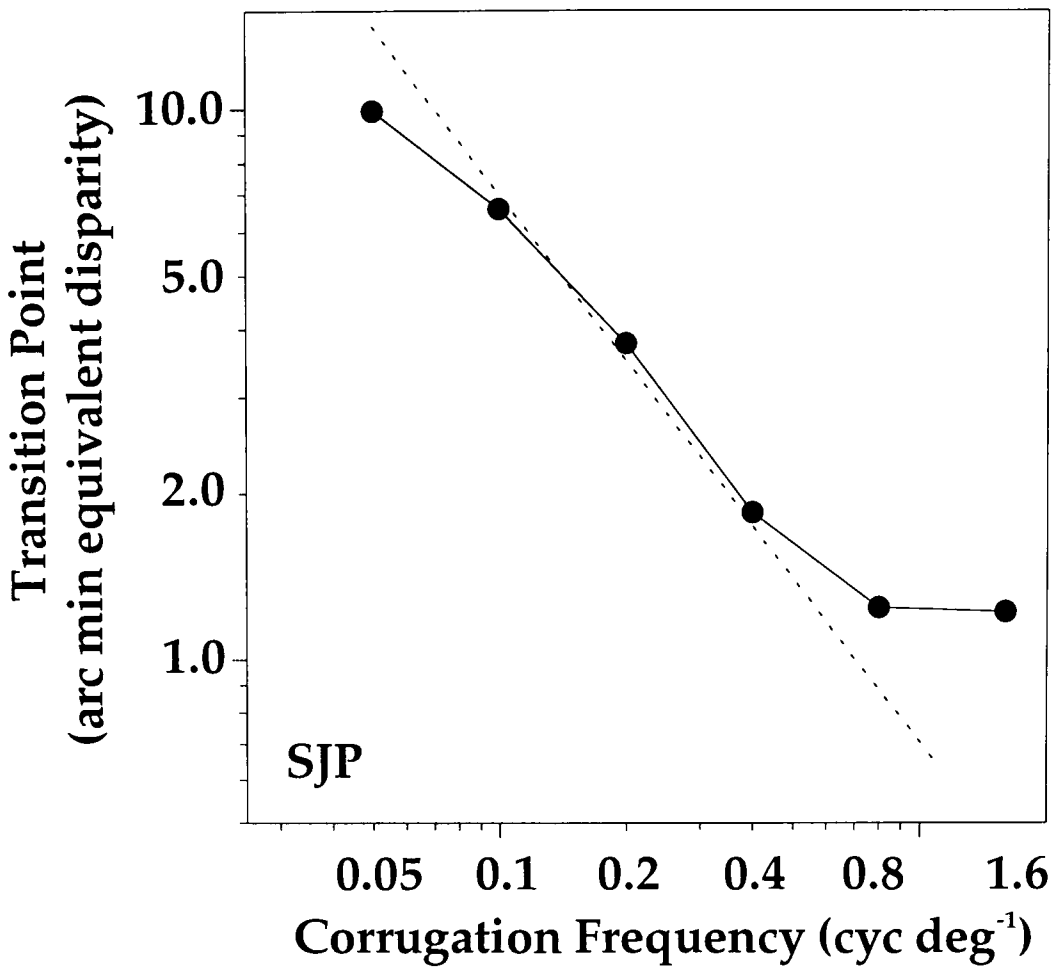


Fig. 3.7: Transition Point as a function of Corrugation Frequency for Subject SJP. The dotted line has a slope of -1.

3.5.4 Discussion

The data are consistent with the notion that over a four-octave range (three-octave for one subject) the Transition Point is principally determined by the MMG present in the stimulus. Since the data do not lie exactly on the line of slope -1 but rather describe slightly shallower functions, it is evident that the MMG is not the sole factor which determines the Transition Points. This issue is addressed in Experiment 1b.

The fact that the function became less steep and then levelled out at corrugation frequencies greater than 0.4 cpd means that above this frequency the Transition Point is dependent on the absolute amount of relative motion. Subjects SAS and SJP required a minimum of about 1 arc min equivalent disparity in order to perceive the stimulus as rotating; subject AMJ required at least 2 arc mins. It will be noted that these values are fairly close to the depth thresholds of these corrugation frequencies. Rogers and Graham (1982) found depth thresholds of about 50 - 70 and 100 arc sec equivalent disparity for corrugation frequencies of 0.8 and

1.6 cpd respectively, using the method of adjustment. Lower thresholds have been found in studies which have used forced-choice methods (Bradshaw and Rogers, 1993). The proximity of the depth thresholds and the Transition Points for these higher corrugation frequencies suggests that such stimuli are perceived to be rotating almost as soon as they become visible, in turn suggesting that the range of relative motion amplitudes over which a stimulus is perceived to be rigid and rotating varies with corrugation frequency. Experiment 5 examines this suggestion in detail.

There is a slight discrepancy between the absolute values of the Transition Points found in Experiment 1a and the results of Rogers and Collett (1989). The latter study reported that a monocularly-viewed parallax surface did not appear to rotate at relative motion values up to 15 arc min equivalent disparity — ".. the perceived surface did not appear either to *deform* or to *rotate* as it translated to and fro." (Rogers and Collett, 1989, p. 702). The highest Transition Point in the present experiment was 11.2 arc min equivalent disparity, meaning that no surface with a greater relative motion value was perceived as stationary. The corrugation frequency of the surface producing this Transition Point was 0.05 cpd, whilst the corrugation frequency of the surface used by Rogers and Collett (1989) was 0.2 cpd. Transition Points for surfaces of this corrugation frequency in Experiment 1a were 3.3 - 4.6 arc min equivalent disparity. The factors accounting for this sizeable difference between data of identical corrugation frequency in the two studies are unclear, but there are two obvious differences between the experimental set-ups, either of which may be responsible.

The effect of oscillation frequency

Firstly, the period of oscillation of the relative motion in the study by Rogers and Collett (1989) was 'about 3 seconds' (p. 700) whilst in Experiment 1a it was about 1 second. Ono and Ujike (1993) and Ichikawa and Ono (1996) have shown that motion thresholds decrease with increasing head/display oscillation frequency. In the former study observers were asked to decrease the amount of relative motion in head-linked parallax displays until motion could no longer be perceived in the stimulus. The head velocity took one of six values between 0.125 and 4 cm

sec^{-1}. It was found that thresholds decreased with a slope of unity as head velocity increased from 0.125 to 0.5 cm sec^{-1}: a decreasing amount of relative motion was required for motion to be perceived in the stimulus as head velocity increased.

Ichikawa and Ono (1996) repeated this experiment under slightly different conditions, using head velocities between 0.5 and 16 cm sec^{-1}. They found that over the whole range of velocities used, the motion threshold decreased with increasing head velocity. Thresholds were also determined for a condition in which the observer remained stationary whilst the whole stimulus moved from side to side. The same pattern of results was obtained.

If the 'motion' perceived by the observers is likened, for the purposes of the current discussion, to the rotation perceived in Experiment 1a, then the data of Ono and Ujike (1993) and Ichikawa and Ono (1996) suggest that the Transition Point may fall when either the head or the stimulus velocity is increased. The oscillation frequencies of Experiment 1a and Rogers and Collett's (1989) study were 1 and 0.33 Hz respectively. If the head/monitor motions are treated as approximate instances of Simple Harmonic Motion (SHM) then these oscillations involve a range of velocities, the maximum values of which are 41 and 14 cm sec^{-1} respectively. Despite the fact that the maximum speed tested by Ichikawa and Ono (1996) was lower than 41 cm sec^{-1}, the graphs of their second experiment suggest that the thresholds would continue to fall at higher velocities.

The difference in oscillation frequency could therefore partly account for the difference between the data of Experiment 1a and that of Rogers and Collett's (1989) experiment.

<u>The effect of method of parallax generation</u>

The second factor which might account for the difference between the results of the two studies concerns the methods used to generate the relative motion. In the study of Rogers and Collett (1989) the display oscilloscope swung to and fro in the frontal plane and the relative motion within the display occurred in temporal synchrony with the oscillations.

Observers rested their heads on a chin-rest and tracked the translating surface. In the current experiment, the relative motion was generated by the side-to-side head movements of the observer.

Rotation is more frequently perceived in displays viewed by stationary observers than in displays viewed by translating observers (Ono and Steinbach, 1990), and motion thresholds are generally lower in the former case (Ichikawa and Ono, 1996). These studies suggest that the propensity to perceive rotation is greater when observers are stationary than when they are moving. It is therefore unlikely that the difference between the parallax generation methods of Experiment 1a and in the study of Rogers and Collett (1989) is the cause of the observed discrepancy between the results of those studies, since this would predict that the discrepancy would lie in the direction opposite to that observed.

The reason for the lack of perception of rotation in the monocular motion parallax condition of Rogers and Collett's (1989) study remains unclear.

The main thrust of the data of the current experiment, however, is that over a four-octave range of frequencies the Transition Point is dependent upon corrugation frequency. Beyond this range it is dependent upon the absolute amount of relative motion.

3.6 Experiment 1b: Transition Point as a function of Corrugation Frequency of Peak-clipped surfaces

3.6.1 Introduction

Whilst the manipulation of corrugation frequency obviously affects the MMG, it also leads to a change in the spacing and the number of stimulus peaks. In order to further tease apart the roles played by these concomitant variations, the present experiment examines Transition Points using peak-clipped sine wave stimuli. This allows a comparison between the Transition Points of stimuli having identical corrugation frequencies but different MMGs.

If Transition Points are dependent on the MMG in the surface, then for a given corrugation frequency, a doubling of the MMG should lead to a halving of the Transition Point. Fig. 3.8 illustrates this prediction. The crosses represent the predicted data from surfaces with MMGs which, per unit Peak-to-Trough amplitude, are twice those of the surfaces represented by the circles. The doubling of MMG is achieved through doubling the amplitude of the surface and 'clipping' the peaks and troughs of the resultant surface at 50% of their respective heights (see Section 3.6.2 below).

For a given corrugation frequency the Transition Point of the clipped surface (Point B) is half that of the unclipped surface (Point A). The figure also illustrates that the Transition Point of the clipped surface (Point B) has the same Transition Point as a surface of twice the corrugation frequency (Point C).

Experiment 1a also found that the function ceased to be dependent on corrugation frequency at a value of 0.8 cpd for two of the subjects, and 0.4 cpd for the other. If this 'cut-off' point is dependent upon the MMG of the surface, then it would be expected that doubling the gradient should lead to a halving of the cut-off corrugation frequency. The cut-off frequencies of the clipped and unclipped surfaces are marked by arrows D and E respectively.

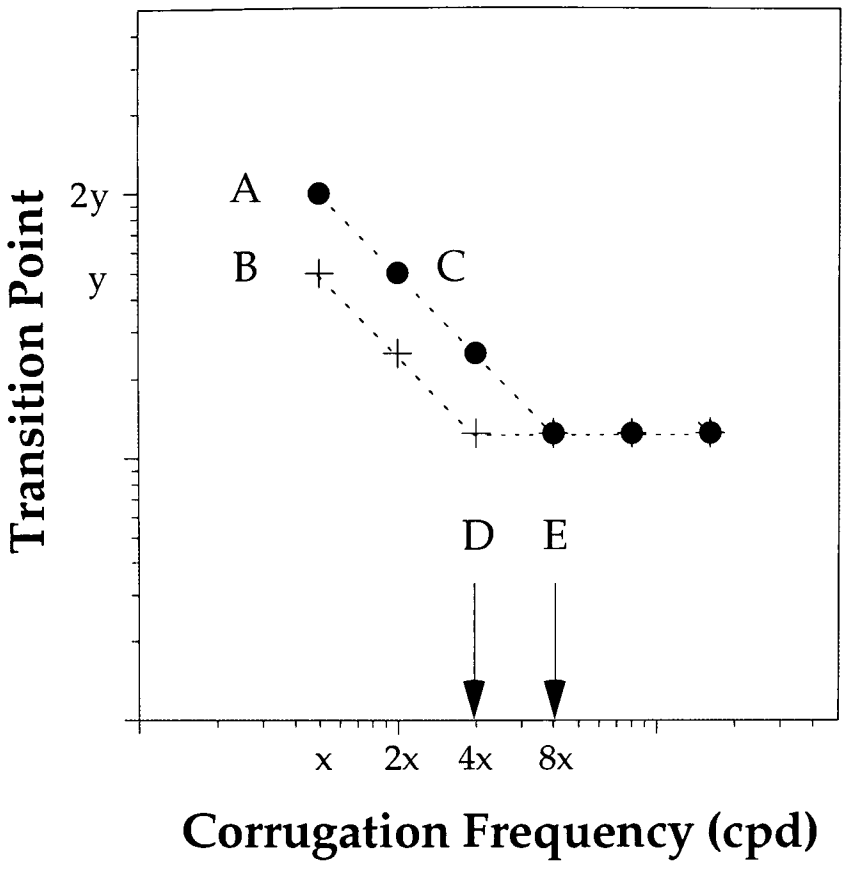


Fig. 3.8: Predicted relationship between Corrugation Frequency and Transition Point for two surfaces. The crosses represent predicted data from surfaces with MMGs which, per unit Peak-to-Trough amplitude, are twice those of the surfaces represented by the circles. See text for details.

The current experiment will thus test three related predictions of the hypothesis that Transition Points are determined by the MMG:

1. The Transition Point will be halved for each doubling of MMG
2. The Transition Points of surfaces with identical MMGs will be identical
3. The 'cut-off' point will halve in frequency for each doubling of the MMG

3.6.2 Stimuli

Stimuli were identical to those of Experiment 1a, except that the sinusoidal depth modulations were clipped. In order to create a clipped stimulus of a given corrugation frequency, the amplitude of the unclipped waveform was doubled and the peaks and troughs of the resulting waveform were clipped at 50% of their respective heights. This produced a stimulus which had the same corrugation frequency but a different MMG as its unclipped original. For a given peak-to-trough amplitude of relative motion, the 50%-clipped waveforms contained MMGs which were double those of the unclipped versions (see Fig. 3.9). 25%-clipped surfaces were

created by quadrupling the amplitude of the unclipped waveform and clipping the resultant peaks and troughs at 25% of their respective heights. For a given peak-to-trough amplitude of relative motion, the 25%-clipped waveforms contained MMGs which were four times those of the unclipped versions.

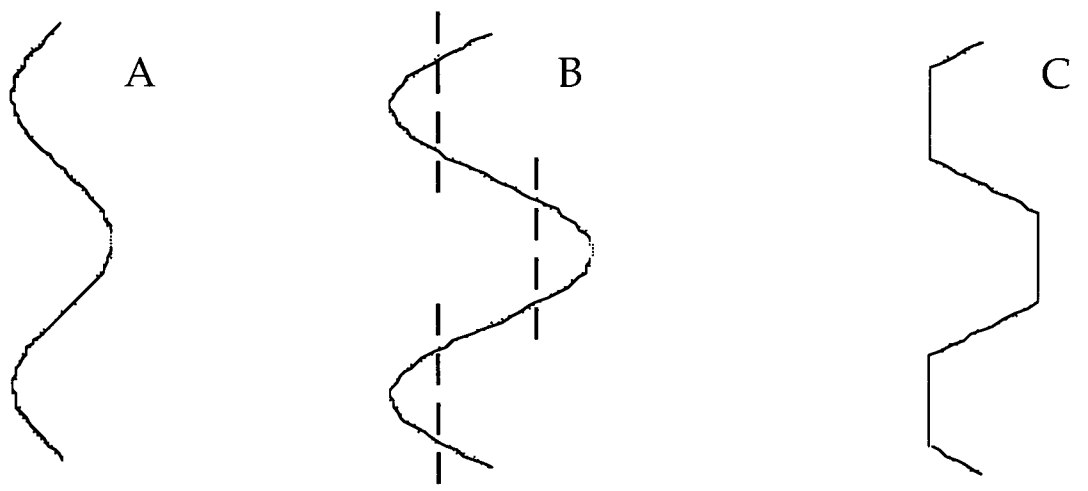


Fig. 3.9: A - Profile of an unclipped surface; B - Profile of an unclipped surface with the same corrugation frequency and twice the peak-to-trough amplitude of A. The dashed lines indicate where clipping will occur; C - Profile of the clipped surface. The maximum motion gradient of C is twice the maximum motion gradient of A. Both stimuli have the same peak-to-trough relative motion amplitude.

3.6.3 Results

Figs 3.10 - 3.12 show Transition Points as a function of corrugation frequency for unclipped, 50%-clipped and 25%-clipped stimuli. Predictions 1 - 3 are considered in turn.

The effect of clipping on the Transition Point

Prediction 1 stated that the Transition Point should be halved for each doubling of MMG. For a given corrugation frequency within the range 0.05 - 0.4 cpd, Transition Points of clipped stimuli were lower than those of unclipped stimuli. For corrugation frequencies up to and including 0.2 cpd, the 50%-clipped waveforms resulted in Transition Points which were on average 64% of those of the unclipped waveforms of the same corrugation frequency. The 25%-clipped waveforms resulted in Transition Points which were on average 40% of those of the unclipped stimuli. The expected percentages, had the Transition Points of all stimuli been determined solely by the MMG, are 50% and 25% respectively. The Transition Points of stimuli of corrugation frequency 0.4 cpd and greater

were only slightly affected by peak-clipping: clipped and unclipped stimuli of similar corrugation frequency had similar Transition Points.

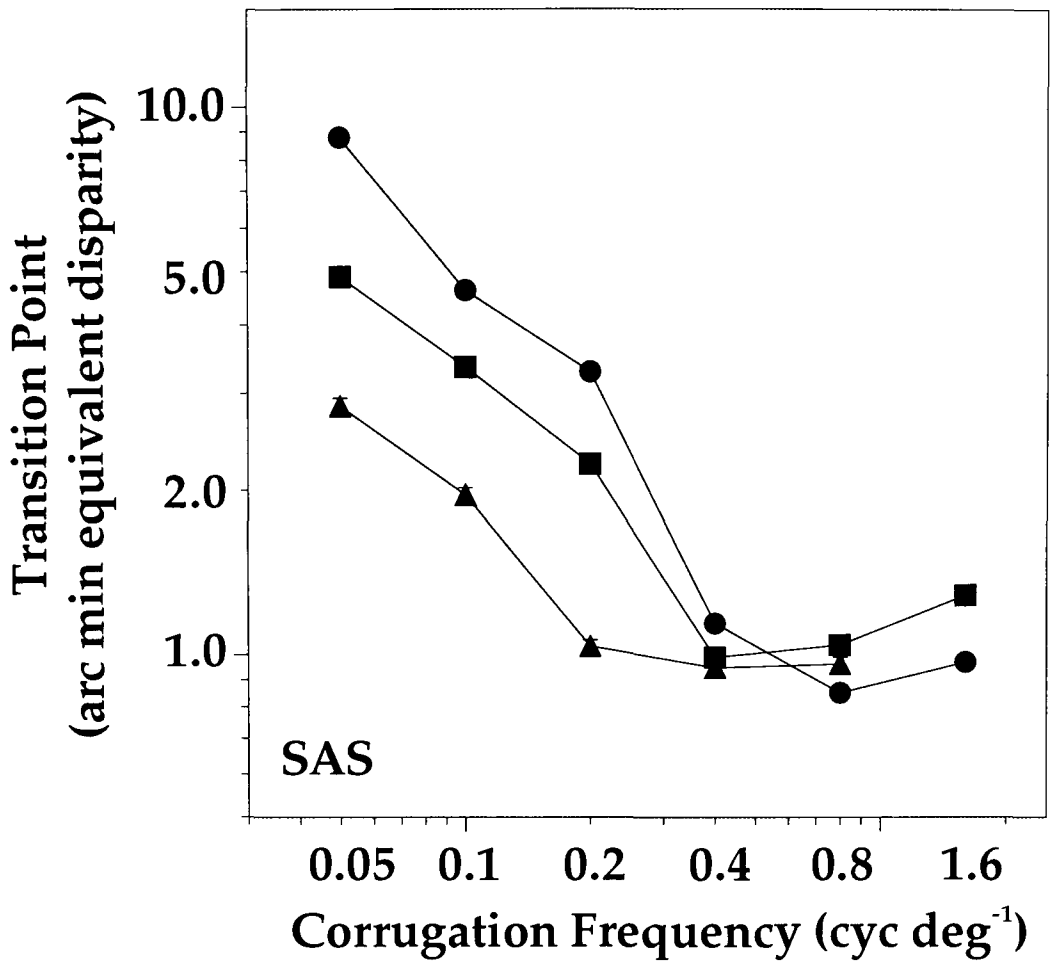


Fig. 3.10: Transition Point as a function of Corrugation Frequency for Subject SAS. The filled circles represent data for the unclipped waveforms, and the filled squares and triangles represent data for the 50%- and 25%-clipped waveforms respectively.

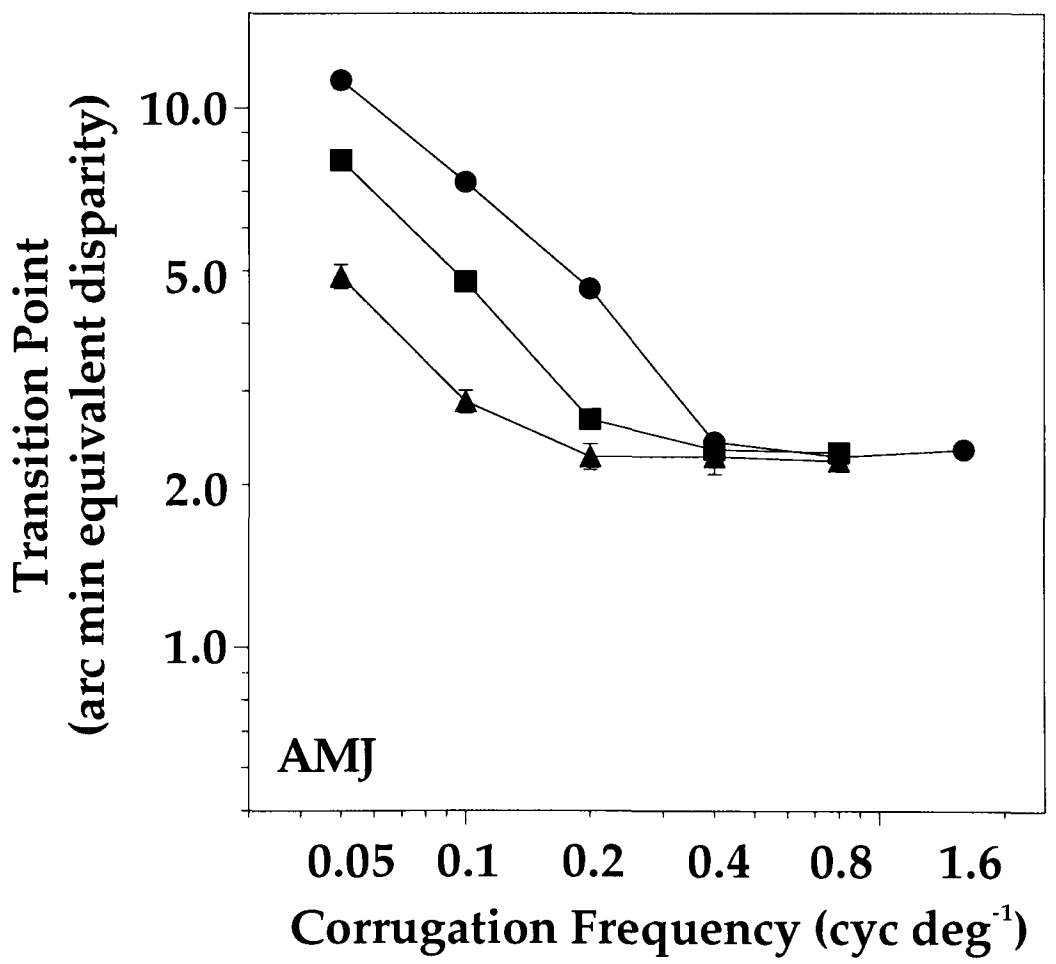


Fig. 3.11: Transition Point as a function of Corrugation Frequency for Subject AMJ. The filled circles represent data for the unclipped waveforms, and the filled squares and triangles represent data for the 50%- and 25%-clipped waveforms respectively.

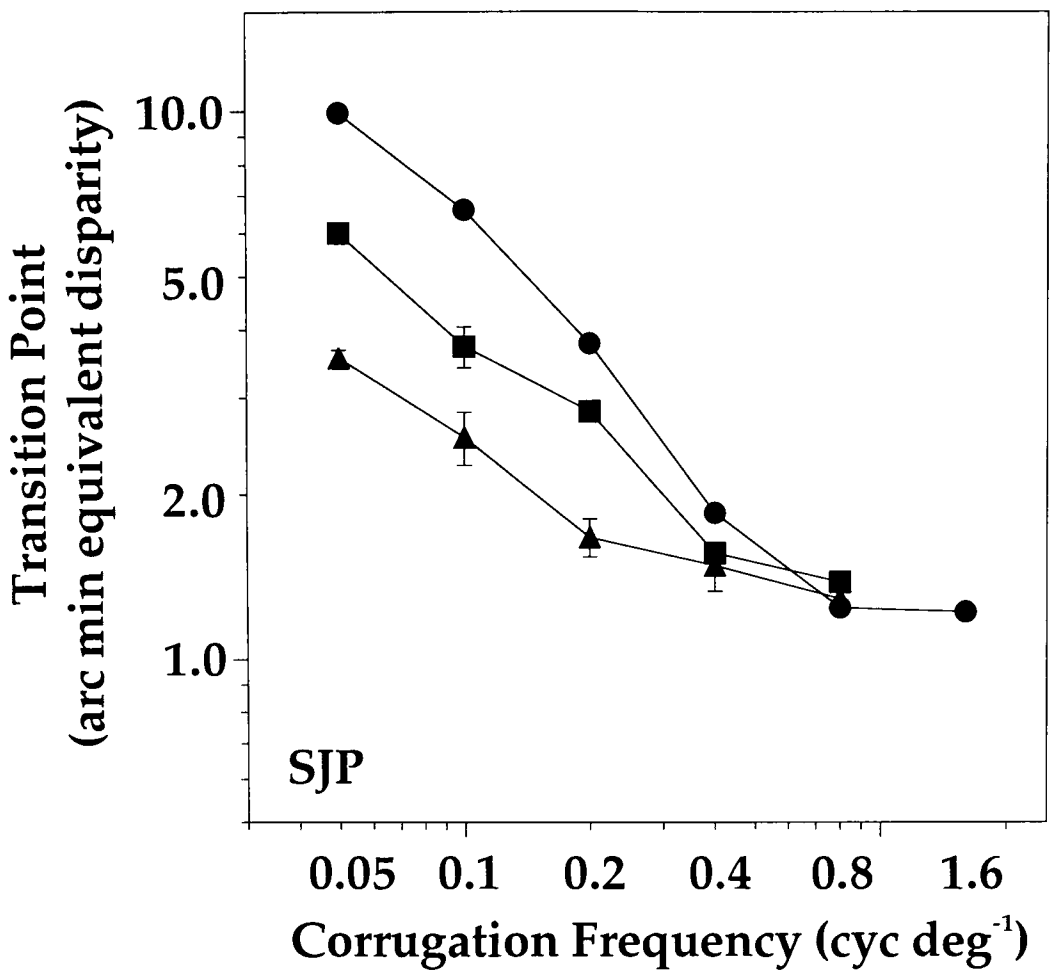


Fig. 3.12: Transition Point as a function of Corrugation Frequency for Subject SJP. The filled circles represent data for the unclipped waveforms, and the filled squares and triangles represent data for the 50%- and 25%-clipped waveforms respectively.

Transition Points of surfaces with identical MMGs

Prediction 2 stated that the Transition Points of surfaces with identical MMGs should be identical. In order to examine the effect of the MMG on the Transition Point more closely, the notion of the disparity gradient (Burt and Julesz, 1980) was applied to the data. The disparity gradient between two stimuli is the difference in disparity between the stimuli, divided by their binocular separation. The disparity gradient between a pair of points lying along the vertical meridian relates their horizontal disparity to their vertical angular separation. The equivalent disparity gradient between two points on a parallax surface (for instance a 'peak' dot and a 'trough' dot) can be defined as the peak-to-trough amplitude, measured in arc min equivalent disparity (see Fig. 1.3), divided by the vertical angular separation of the dots, measured in arc min. Note that this assumes the surface between the two dots has a constant motion gradient. In the case of sinusoidal depth variations in the surface, the MMG is calculated as follows.

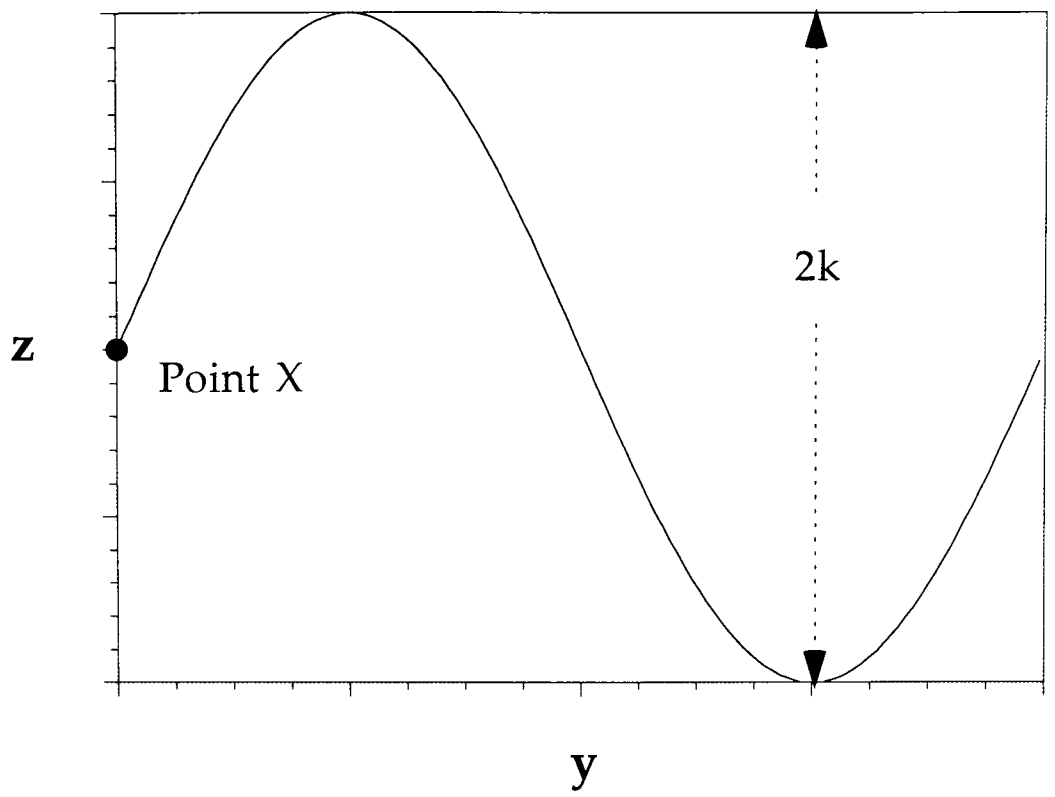


Fig. 3.13: Representation of a horizontally-oriented, sinusoidally corrugated surface of corrugation frequency f and peak-to-trough amplitude 2k. The gradient of the function is steepest at Point X.

Fig. 3.13 represents a horizontally-oriented, sinusoidally corrugated surface of corrugation frequency f and peak-to-trough amplitude 2k.

The equation of the surface is

$z = k\sin(\omega y)$, where

$\omega = 2\pi f$.

The gradient of the surface is given by

$dz/dy = k\omega\cos(\omega y)$.

Point X represents the point at which the function is steepest. Thus the MMG is given by

$dz/dy = k2\pi f$.

Figs. 3.10 - 3.12 show that at relatively low values of the MMG the Transition Points of surfaces with identical MMGs were similar. The

Transition Points of surfaces with MMGs at or beyond the cut-off point were also similar.

The data are replotted in Figs. 3.14 - 3.16 below to show the effect of the MMG more clearly. Transition Points are plotted as a function of MMG per unit Peak-to-Trough amplitude. Filled circles, open squares and crosses represent the Transition Points of unclipped, 50%- and 25%-clipped surfaces respectively.

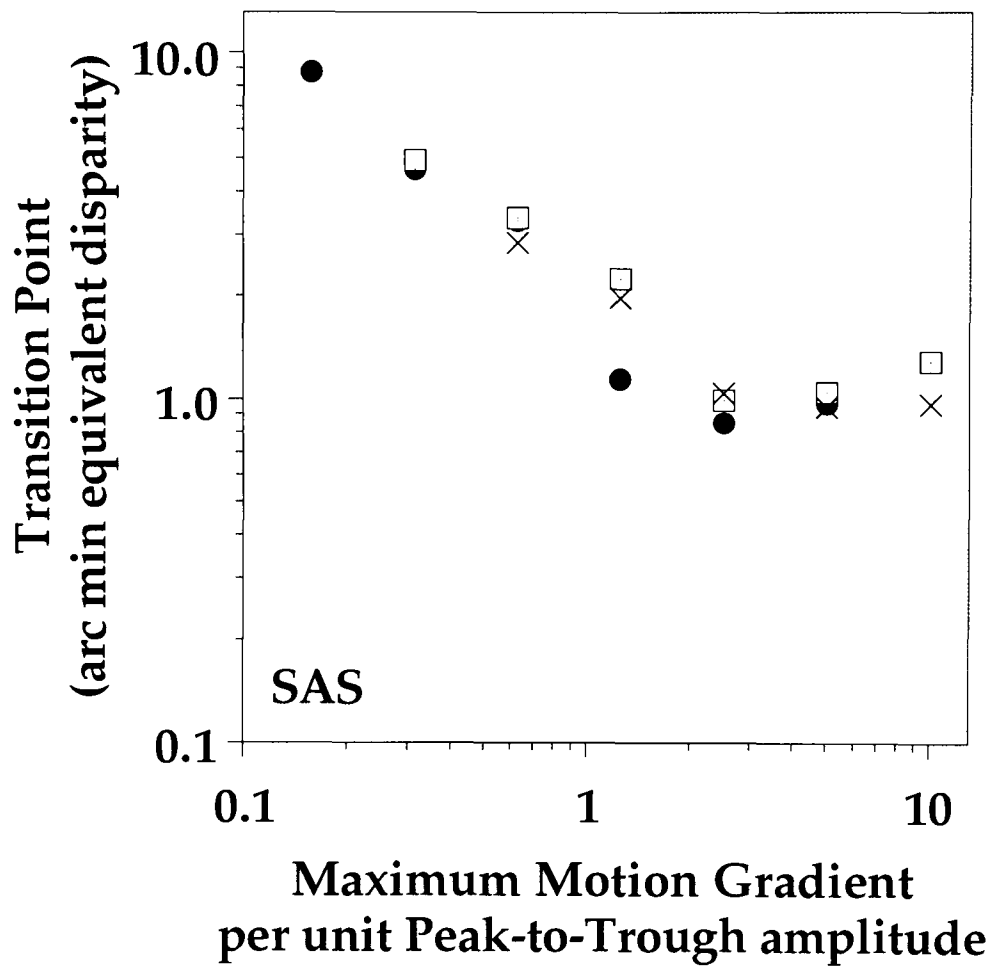


Fig. 3.14: Transition Point as a function of MMG per unit Peak-to-Trough amplitude for Subject SAS.

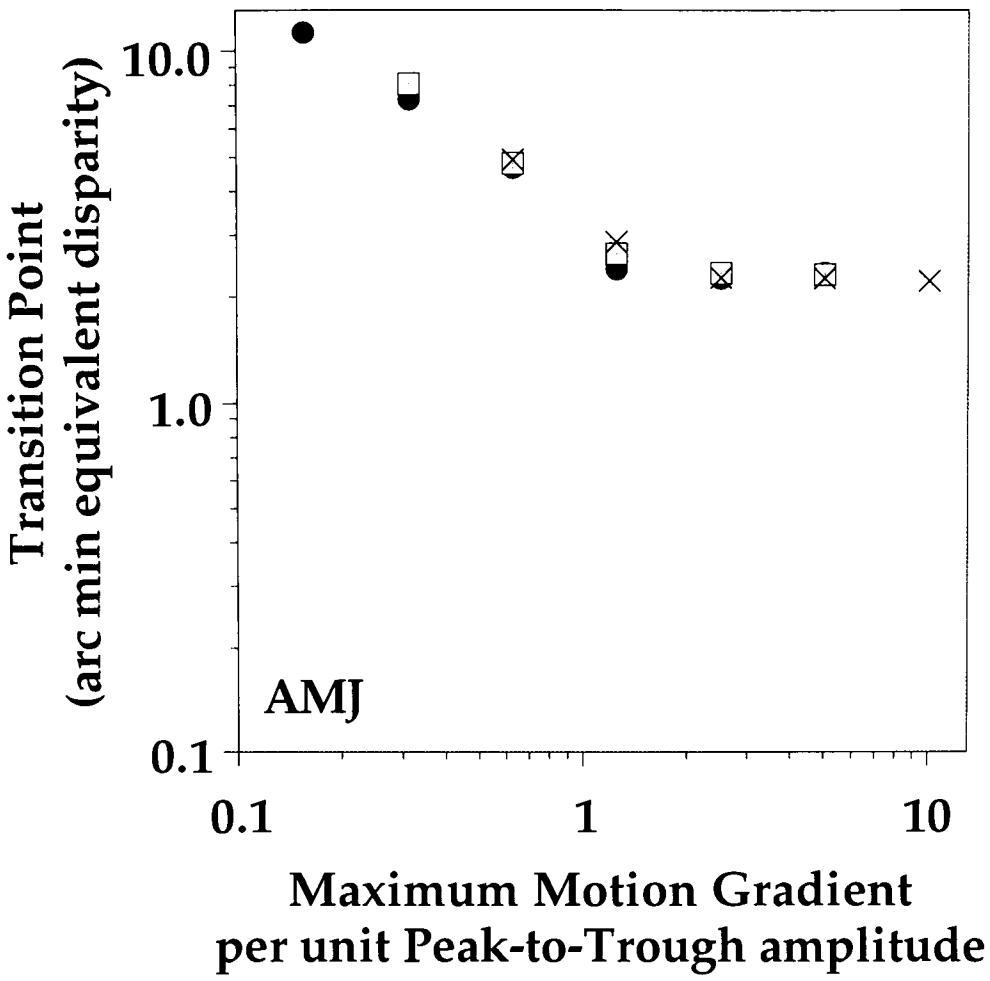


Fig. 3.15: Transition Point as a function of MMG per unit Peak-to-Trough amplitude for Subject AMJ.

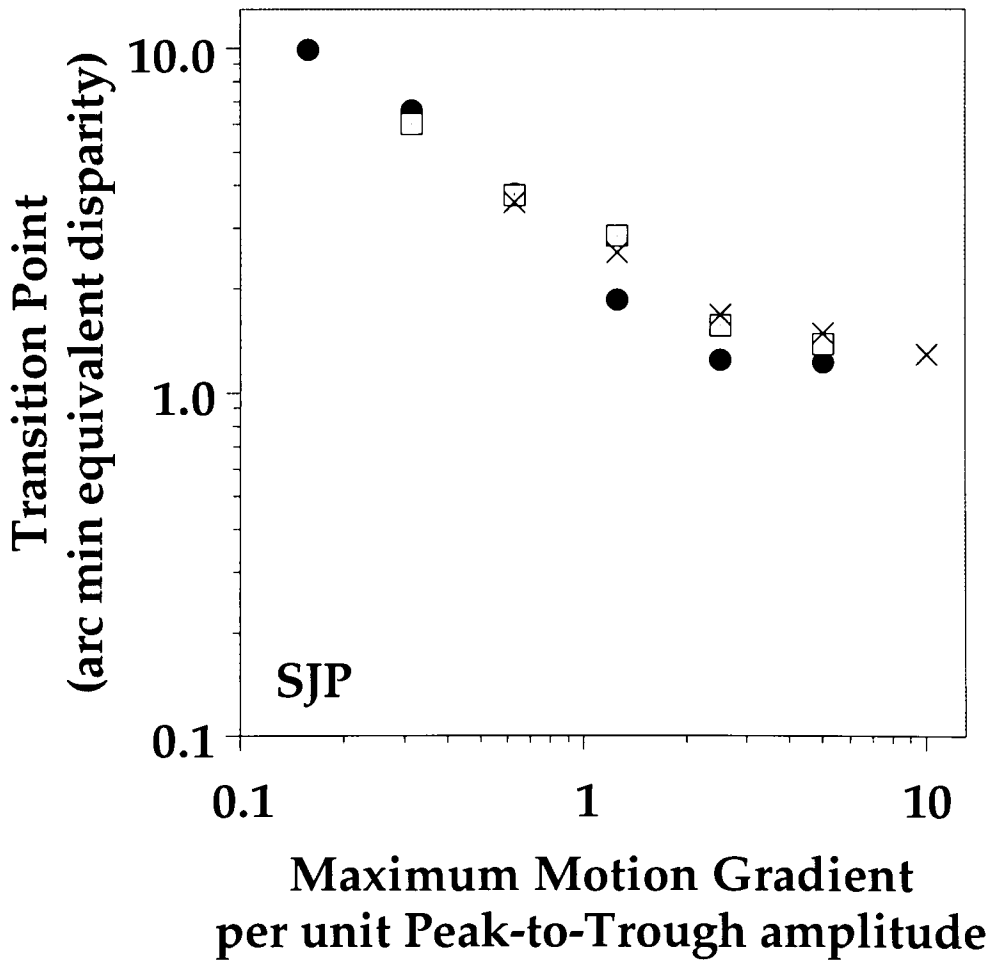


Fig. 3.16: Transition Point as a function of MMG per unit Peak-to-Trough amplitude for Subject SJP.

Prediction 2, namely that Transition Points of surfaces of a given MMG per unit Peak-to-Trough amplitude should be similar, was partly borne out. It was largely satisfied for all clipped surfaces: the separation of Transition Points of clipped surfaces with identical MMGs was always less than one arc min equivalent disparity. For Subjects SAS and SJP, the data from the unclipped surfaces sometimes fell below those for the clipped surfaces. For Subject AMJ the prediction was conclusively borne out.

The effect of clipping on the cut-off point

Prediction 3 stated that the corrugation frequency of the cut-off point will halve for each doubling of the MMG. In order to obtain the theoretical cut-off points, two power functions of the form $y = ax^b$ were fit to each set of data. Power functions of this nature produce linear fits on log-log axes.

The method for fitting these power functions was as follows. Each possible pair of power functions was fitted to each set of data (for instance, for the unclipped data, best-fitting lines for the following corrugation frequency ranges were fitted:

1. 0.05 - 0.1 cpd; 0.1 - 1.6 cpd
2. 0.05 - 0.2 cpd; 0.2 - 1.6 cpd
3. 0.05 - 0.4 cpd; 0.4 - 1.6 cpd
4. 0.05 - 0.8 cpd; 0.8 - 1.6 cpd).

The two R-values associated with each pair of fits were combined in a weighted fashion such that the contribution of each R-value was weighted by the number of data points in the corresponding line of best fit. Thus for example in the first pair of power functions above, the R-values for the lines fitted to the 0.05 - 0.1 cpd data and the 0.1 - 1.6 cpd data respectively contributed 2/7 and 5/7 of the combined R-value.

The pair of fits which resulted in the greatest combined R-value was plotted for each set of data. The intersection of these functions was taken as the cut-off point. On Figs. 3.17 - 3.19 the cut-off points are marked by arrows. The closed, open and line arrows represent the cut-off points of the 25%-, 50%- and unclipped stimuli respectively. The equations for the steeper portion of each function are shown. The equations for the

shallower portions are in Appendix B.

The cut-off points are approximately in line with Prediction 3. The graphs also show that, for all subjects, the gradients of the 25%-clipped functions are shallower than those of the unclipped functions.

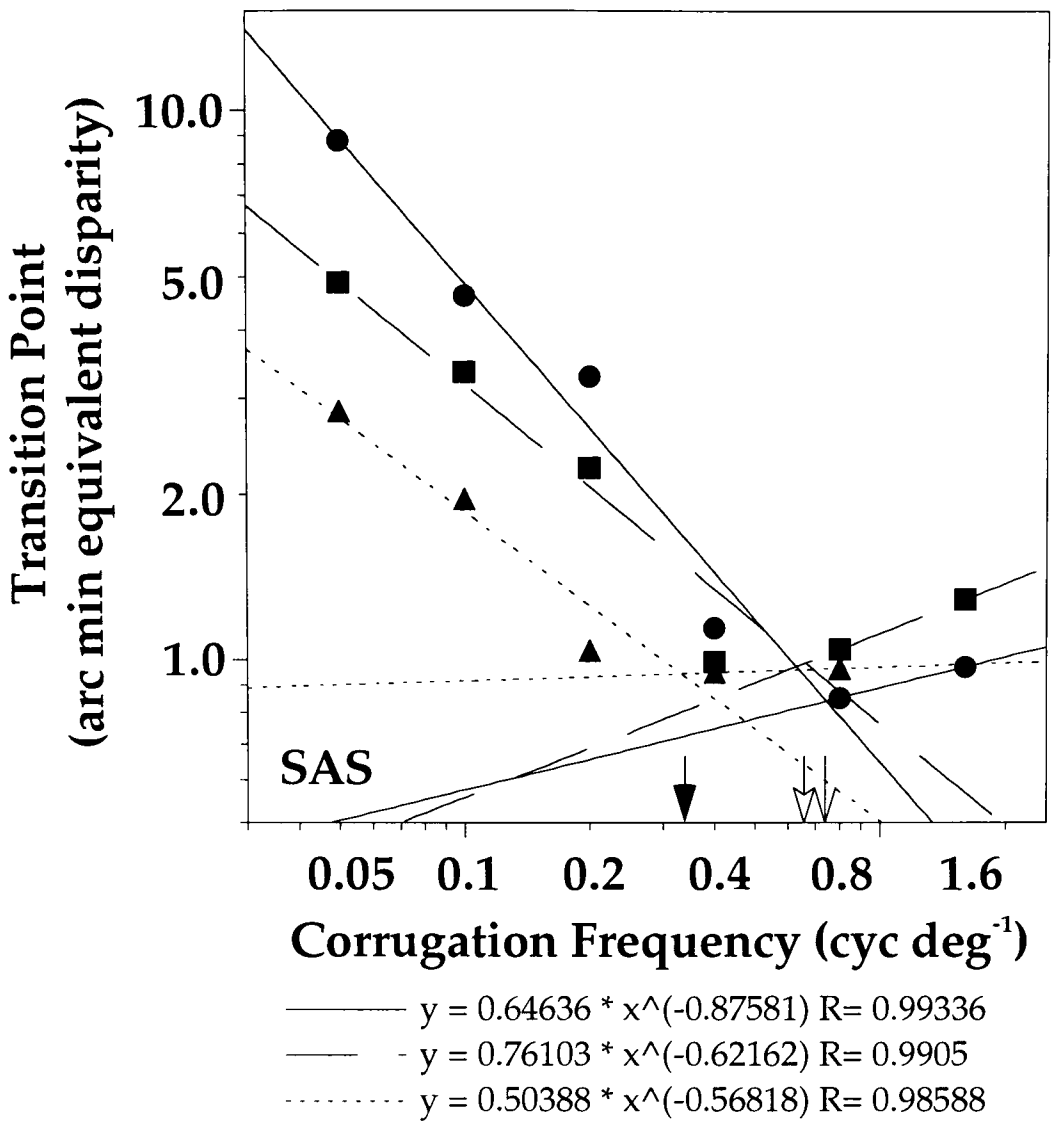


Fig. 3.17: Transition Point as a function of Corrugation Frequency, for unclipped (circles), 50%- (squares) and 25%- (triangles) clipped surfaces, for Subject SAS. Best-fitting lines of the form y = ax^b are shown. The arrows, from right to left, represent theoretical cut-off points for unclipped, 50%- and 25%-clipped surfaces respectively.

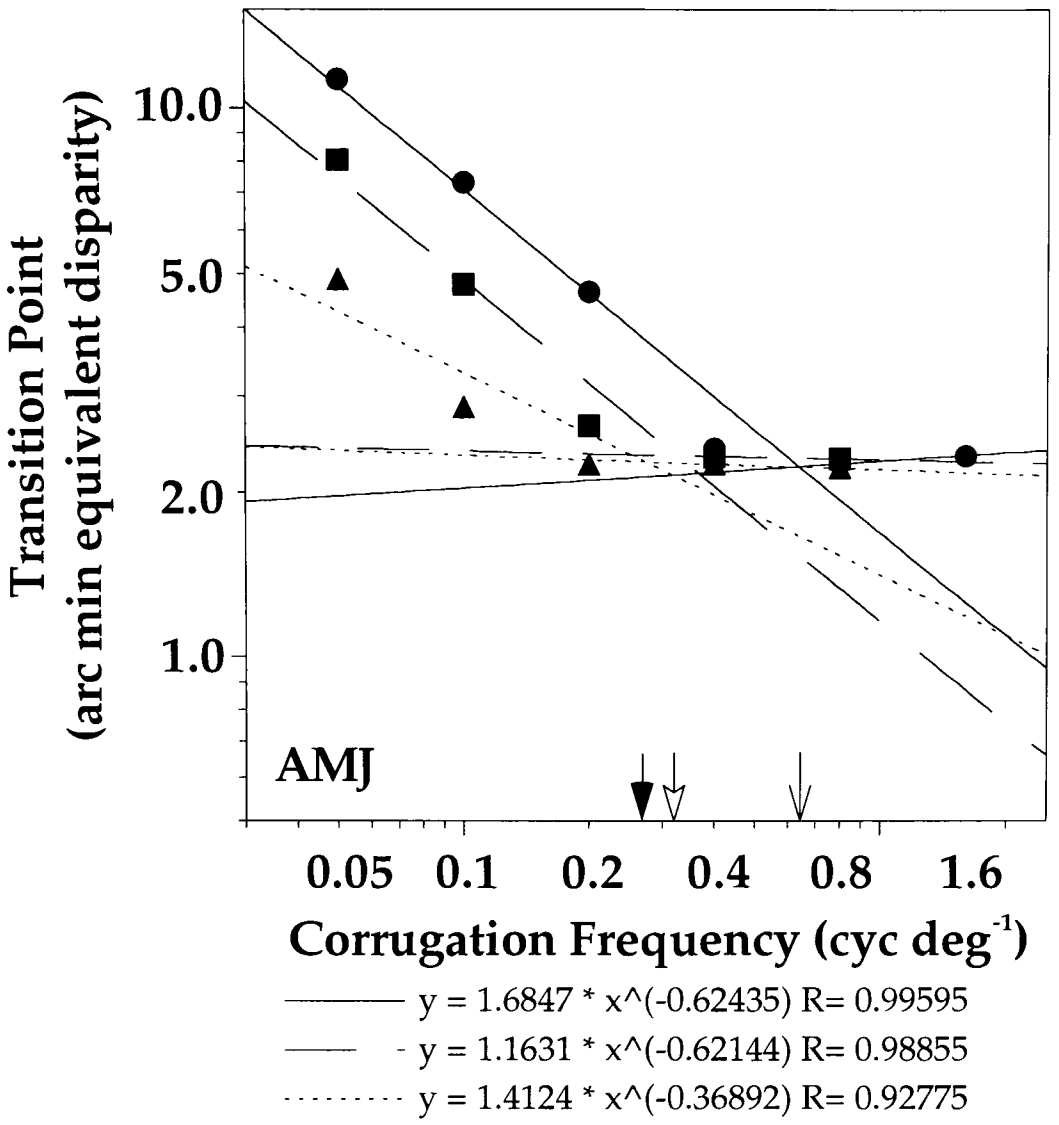


Fig. 3.18: Transition Point as a function of Corrugation Frequency, for unclipped (circles), 50%- (squares) and 25%- (triangles) clipped surfaces, for Subject AMJ. Best-fitting lines of the form y = ax^b are shown. The arrows, from right to left, represent theoretical cut-off points for unclipped, 50%- and 25%-clipped surfaces respectively.

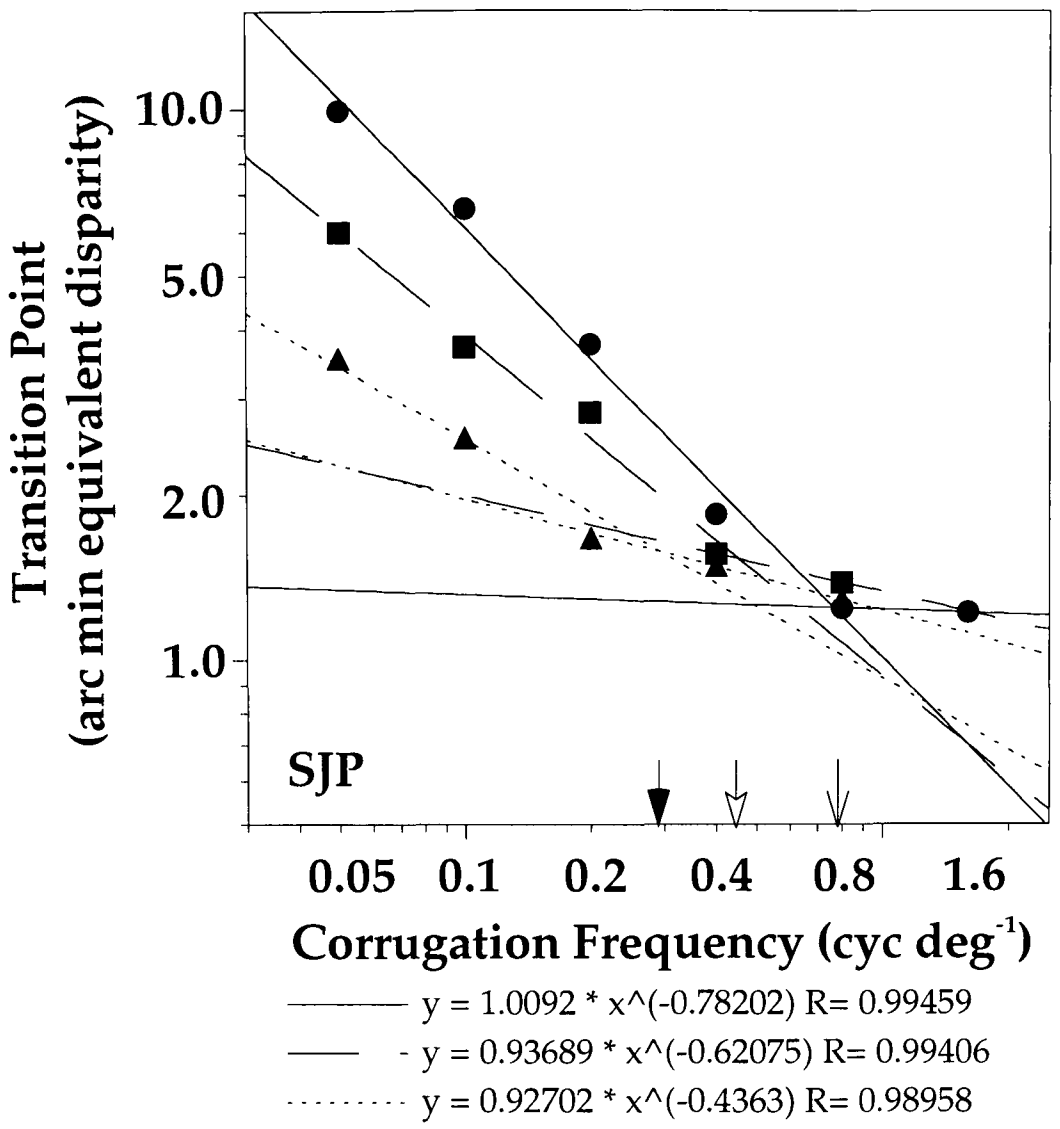


Fig. 3.19: Transition Point as a function of Corrugation Frequency, for unclipped (circles), 50%- (squares) and 25%- (triangles) clipped surfaces, for Subject SJP. Best-fitting lines of the form y = ax^b are shown. The arrows, from right to left, represent theoretical cut-off points for unclipped, 50%- and 25%-clipped surfaces respectively.

3.6.4 Discussion

The Transition Points of clipped surfaces were generally lower than those of unclipped surfaces with the same corrugation frequency. Transition Points of the 25%-clipped surfaces were lower than those of the 50%-clipped surfaces. This is qualitatively in line with the hypothesis that Transition Points are determined by the MMGs. The Transition Points of the clipped surfaces relative to the unclipped surfaces were slightly greater than predicted, however.

Figs. 3.14 - 3.16 illustrate that surfaces with identical MMGs per unit amplitude have similar Transition Points. This provides further evidence that, for a range of MMGs per unit Peak-to-Trough amplitude, Transition Points are determined by the MMG. At the higher levels of MMG per unit Peak-to-Trough amplitude (those beyond the cut-off point), the similarity

of Transition Points is due to the increasing dependence upon the absolute relative motion amplitude.

Figs. 3.17 - 3.19 show that the gradients of the functions are slightly less than the predicted value of -1. The gradients are generally greatest for the unclipped surfaces and lowest for the 25%-clipped surfaces, indicating that the MMG plays a decreasing role in the determination of the Transition Point as the degree of surface 'clippedness' increases. The Transition Point is primarily determined by the MMG for surfaces containing MMGs less than or equal to the MMG at the cut-off Points.

Under some conditions the Transition Points of the clipped waveforms are higher than those of unclipped waveforms of equal maximum motion gradient. For example, the Transition Points of 50%-clipped stimuli of corrugation frequency 0.2 cpd are greater than those of unclipped 0.4 cpd stimuli for subjects SAS and SJP. This suggests that even within the range where Transition Points are primarily determined by the motion gradient, they are affected by other factors also. It is possible that peak-spacing is one of these factors: lower corrugation frequencies contain both lower maximum motion gradients *and* greater peak spacing, and it may be the case that the higher Transition Points of these stimuli are due to both factors. Another possible factor is the form of the stimulus. Clipped sine wave stimuli obviously contain plateaux of constant relative motion at their peaks and troughs, and it is perhaps the presence of these regions that causes Transition Points to be higher than expected. Ono and Steinbach (1990) found that observers perceived square waves to shear rather than to rotate. Clipped surfaces, due to their plateaux of constant relative motion, bear a considerable resemblance to square waves and thus may contain a shearing component to their motion. It may be the case that a small degree of shearing is perceived at the expense of rotation, thus resulting in the slight elevation of the Transition Point.

To conclude, these data indicate that the MMG is an important determinant of the Transition Point for a range of surfaces. The dependence of the Transition Point on the MMG decreases with increasing corrugation frequency, and with the increasing 'clippedness' of the surface.

3.7 Experiment 2: Transition Point as a function of Waveform Type

3.7.1 Introduction

Ono and Steinbach (1990) presented observers with motion parallax surfaces containing one of four different velocity gradients: sine, triangle, ramp (sawtooth) or square waveform. The relative motion amplitude was either 12 or 113 arc min equivalent disparity. It was found that the amount of perceived motion in the parallax stimuli was dependent upon the waveform. The square waveform was perceived to contain the most motion, followed by the ramp, triangle and sine waveforms respectively. At first sight these results imply that the Transition Points of the waveforms might vary in the reverse order, that is, the sine waveform would have the highest Transition Point, followed by the triangle, ramp and square waveforms respectively. However, as noted previously, the type of motion perceived in sine, triangle and ramp waveforms is generally of a rotary nature, while the perceived motion of square waveforms is not (Ono and Steinbach, 1990). The Transition Point, as defined in Section 3.3, relates to the perceived *rotation* of the surface. Since the square wave surface was not reported as rotating, the results of Ono and Steinbach do not provide clear predictions of the ordering of the Transition Points.

One hypothesis is that the Transition Points of these waveforms are determined by the visual system's responses to the Fourier components of the waveform. Fourier analysis is the process whereby a stimulus is transformed into a weighted sum of harmonic functions. The response of the visual system to the individual components of a compound stimulus has been used to explain the overall response of the system to the stimulus as a whole. This analysis was initially applied to the luminance domain (e.g. Campbell and Robson, 1968). More recent evidence suggests that the same principles may be applied to the stereo domain.

The suggestion that the visual system may analyse disparity within bandpass disparity spatial frequency and disparity orientation channels has been supported by Tyler (1975), who found tilt- and size- aftereffects following the prolonged inspection of cyclopean disparity gratings.

Schumer and Ganz (1979) performed a selective adaptation experiment and showed that the dectectability of a target grating, following prolonged viewing of a high-amplitude disparity grating, was dependent upon its relative (disparity) corrugation frequency. Thresholds were elevated for detection of a target with a similar spatial frequency to that of the adapting stimulus, but became decreasingly elevated as the difference between the spatial frequencies of the gratings increased. Using masking techniques, in which the task is to detect a target grating in the presence of maskers of different spatial frequency properties, Tyler, (1983), Cobo-Lewis and Yeh (1994) and Prince (1998) have shown that the threshold is maximally elevated when the target and masker have similar spatial frequencies, with a falling off in elevation as the difference between the spatial frequencies increases.

Various studies have demonstrated similarities between depth perception through motion parallax and stereopsis, which suggests that these two sources of depth information may be processed in similar ways (e.g. Rogers and Graham, 1982; Graham and Rogers, 1982a, 1982b; Rogers and Graham, 1984; Bradshaw and Rogers, 1996; see Section 1.4). Indeed, using the 'notched noise' technique of Patterson (1976), Hogervorst, Bradshaw and Eagle (1998) obtained results consistent with the notion that the motion parallax sensitivity function comprises a series of independent channels. On these grounds it might be expected that the response of the system to a square wave, for instance, could be predicted by its response to the individual frequency components.

However, there is evidence to suggest that responses to compound motion stimuli are *not* determined by responses to the individual components. Sachtler and Zaidi (1995) measured observers' sensitivity to motion for compound stimuli comprising frequencies 0.133 and 0.4 cpd, or 0.4 and 1.2 cpd. The amplitude of the higher frequency of each pair was always set to one third the amplitude of the lower (note also that the higher frequency of each pair is the third harmonic of the lower frequency). The higher frequency was either added to or subtracted from the lower frequency, leading to profiles approximating square and triangle waves respectively. It was found that sensitivity was dependent on the relative phase of the components. It was greater for the phase combination with the steeper velocity gradient and smaller peak amplitude, even when

the amplitudes of the sinusoidal components were equated for the two compound waveforms. The authors concluded that the mechanisms underlying the detection of motion are particularly sensitive to steep velocity gradients. Physiological evidence to support this claim was provided by Reppas, Sourabh Niyogi, Dale, Sereno and Tootell (1997). Observers were presented with random dot patterns which alternated between periods of segmented and uniform motions. Motion boundaries were created in three ways: by shearing motions, in which local velocities were parallel to the motion boundary, by compressive motions, in which local motion was orthogonal to the boundary, and by a checkerboard pattern which contained both types of edge. A robust functional magnetic resonance imaging (fMRI) signal, which was selective for motion segmentation, was found.

Welch (1996) measured thresholds of discrimination between a sine wave of frequency f_1, and a square wave (which contains frequencies f_1, f_3, f_5,.... with amplitudes decreasing in inverse proportion to frequency). In the vein of Campbell and Robson's (1968) reasoning it was proposed that if the visual system's response to the square wave is based upon analyses of its components then the discrimination threshold should depend upon the detectability of the f_3 component. However, it was found that discrimination thresholds were lower than the f_3 detection thresholds.

A further experiment compared the threshold of discrimination between a compound waveform containing frequencies f_1 and f_3 in peaks-subtract (square wave) phase and a square waveform. If the visual system carried out an analysis of the stimuli which was based on analyses of their individual components then it would be predicted that the discrimination threshold would equal the f_5 threshold. Discrimination thresholds were, however, lower than the f_5 detection threshold. Discrimination thresholds between a triangle waveform and a stimulus combining frequencies f_1 and f_3 in peaks-add (triangle wave) phase were more variable. It was concluded that observers were not responding on the basis of the individual components.

Schumer's (1979) unpublished dissertation (cited by Norman, Lappin and Zucker, 1991) reports similar results to those above. He compared the discriminability of a stereoscopic square wave and its fundamental

(1/3 cpd), with the detectability of the third harmonic of the square wave (1 cpd). Discriminability was higher than could be accounted for by the detectability of the harmonic alone.

Using stereoscopic surfaces, Norman et al. (1991) asked observers to discriminate between a triangle wave and a surface obtained by adding the fundamental and third harmonic of the triangle wave. It was found that discrimination thresholds were lower than predicted by the detectability of the higher frequency components. The authors suggested that the low thresholds were due to the difference in continuity or curvature of the stimuli, rather than to differences in their power spectra.

The present experiment investigated Transition Points as a function of waveform. If the Transition Points of compound stimuli are based upon the Transition Points of the constituent sine waves, then it would be predicted that they would lie beneath those of the sine wave of the same frequency as the compound. If, however, the motion *gradients* are important then it would be predicted that waveforms with similar gradients would have similar Transition Points.

3.7.2 Stimuli

Stimuli were modulated in depth according to one of four types of waveform: sine, triangle, square and ramp. Corrugation frequencies were 0.05, 0.1, 0.2, 0.4, 0.8 and 1.6 cpd. Stimulus size was 20°.

3.7.3 Procedure

This was identical to that of the previous experiments of this chapter, except that subjects completed three experimental blocks instead of four, corresponding to a total of 30 trials at each relative motion amplitude within each condition.

3.7.4 Results

Figs 3.20 - 3.22 show Transition Points as a function of corrugation frequency. Circles, triangles, diamonds and squares represent Transition Points of sine, triangle, ramp and square waveforms respectively.

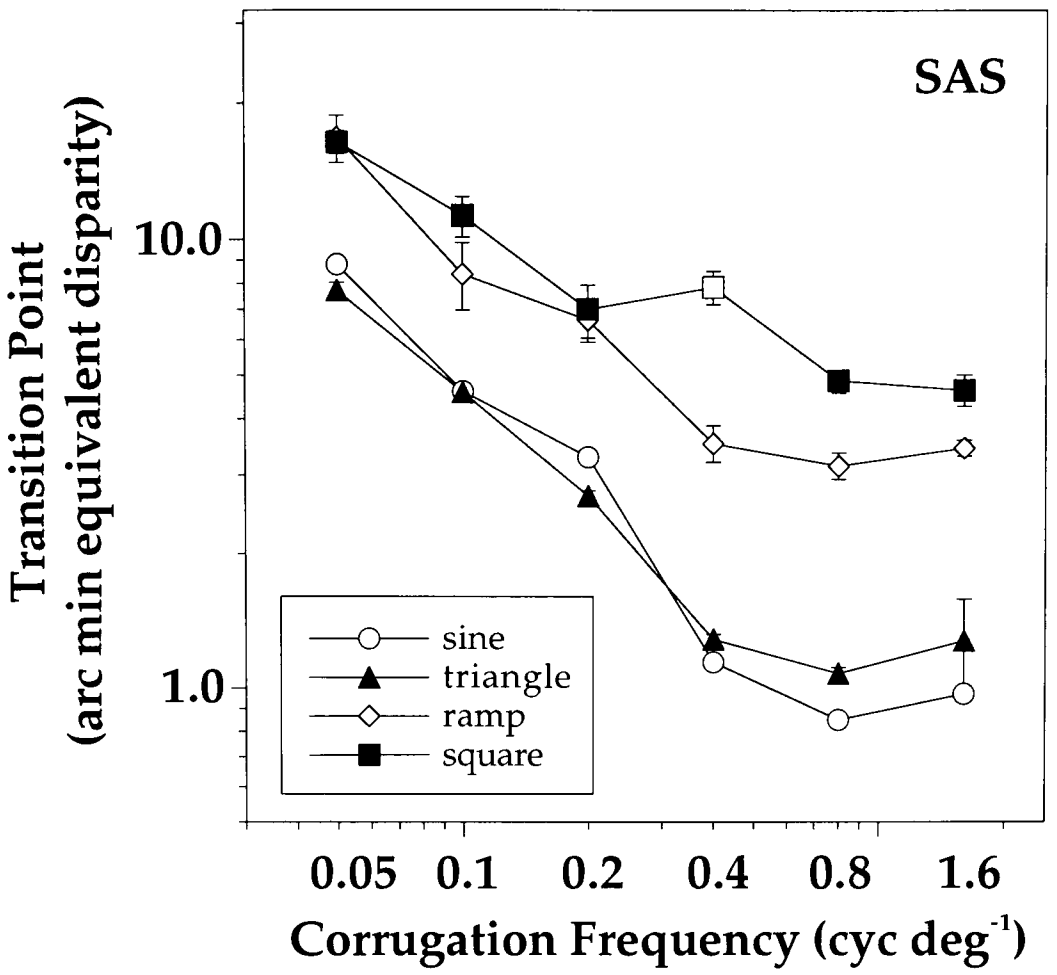


Fig. 3.20: Transition Point as a function of Corrugation Frequency for four waveform types, for Subject SAS.

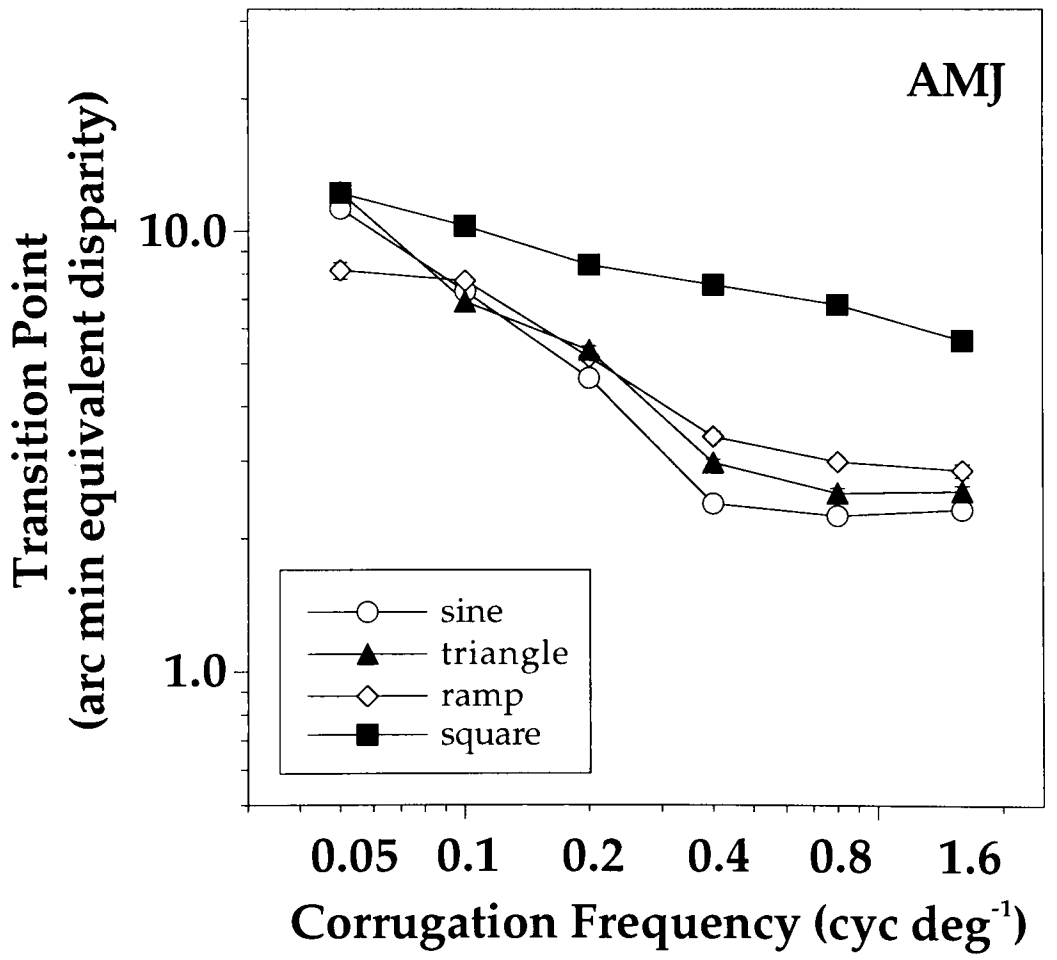


Fig. 3.21: Transition Point as a function of Corrugation Frequency for four waveform types, for Subject AMJ.

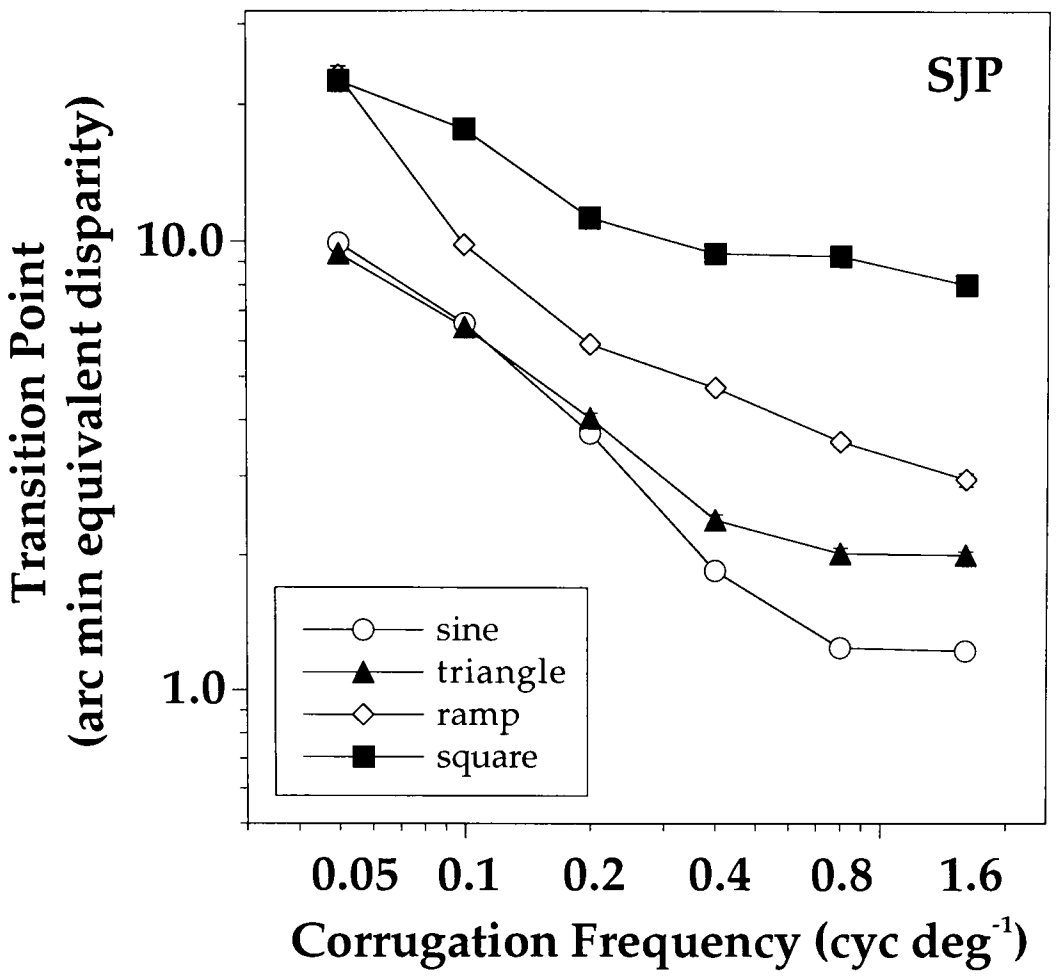


Fig. 3.22: Transition Point as a function of Corrugation Frequency for four waveform types, for Subject SJP.

The data of each subject were entered into a two-way repeated measures analysis of variance (ANOVA). The factors were waveform (4 levels) and corrugation frequency (6 levels).

It will be recalled that experimental trials were blocked in groups of 70, corresponding to ten repetitions of each of seven relative motion amplitudes. In the current experiment, subjects completed three blocks; the combined data were entered into Probit analysis, which produced an estimate of the Transition Point. Raw data of the type required for ANOVAs were unavailable. In order to create the required 'raw data', data from each of the three blocks were entered into three *separate* Probit analysis, thus providing three separate estimates of the Transition Point. These served as the 'raw data' upon which the ANOVA was carried out. It is noted that this is not a perfect method by which to obtain data to enter into the ANOVA — a perfect method does not exist in the current situation, since each Transition Point of Figs. 3.20 - 3.22 does *not* represent the mean of a number of individually measured Transition Points. However, the procedure outlined above was deemed the most suitable.

Significant main effects of waveform were found for every subject (Subject SAS: [F(3, 6), p<0.0005], Subject AMJ: [F(3, 6), p<0.0005], Subject SJP: [F(3, 6), p<0.0005]). There was also a significant main effect of corrugation frequency for every subject (Subject SAS: [F(5, 10), p<0.0005], Subject AMJ: [F(5, 10), p<0.0005], Subject SJP: [F(5, 10), p<0.0005]). There was a significant interaction for every subject (Subject SAS: [F(15, 30), p<0.0005], Subject AMJ: [F(15, 30), p<0.0005], Subject SJP: [F(15, 30), p<0.0005]).

The presence of a main effect of waveform means merely that the Transition Points of at least one waveform differ significantly from the Transition Points of one other waveform. In order to determine exactly which of the waveforms' Transition Points were significantly different, (post-hoc) Tukey HSD tests were carried out. This test fixes the familywise error rate (FW) at α against all null hypotheses; it is regarded as the best procedure for controlling FW when making all pairwise comparisons between means.

The results of the Tukey's HSD tests are found in Appendix C. These show that, for Subject SJP, none of the differences between the Transition Points of the sine and triangle waveforms were significant. For Subject SAS, only the difference between the Transition Points of the sine and triangle waveforms of corrugation frequency 0.05 cpd was significant. Almost all other differences between Transition Points were significant. Exceptions were the non-significant differences between the Transition Points of the ramp and square waveforms of corrugation frequency 0.2 cpd (Subject SAS) and corrugation frequency 0.05 (Subject SJP), and the non-significant difference between the Transition Points of the triangle and ramp waveforms of corrugation frequency 1.6 cpd (Subject SJP).

The results for Subject AMJ were slightly less clear: the differences between the Transition Points of the sine and triangle waveforms were significant for corrugation frequencies of 0.05, 0.2 and 0.8 cpd, and were non-significant for the remaining corrugation frequencies. The differences between the Transition Points of the square waveform and the Transition Points of other waveforms were always significant. There was no clear pattern of significant differences between the Transition Points of the remaining waveforms. However, Fig. 3.21 shows that the *ordering* of the four waveforms' Transition Points for this subject was identical to that of

the remaining subjects, at the three highest corrugation frequencies at least — namely, Transition Points were lowest for the sine waveform and highest for the square waveform. Transition Points for the triangle and ramp waveforms fell between these values.

For subjects SAS and SJP, the Transition Point functions of the ramp and square waveforms lay at significantly higher relative motion values than the functions for the sine and triangle waveforms. For subject AMJ the square waveform function was also significantly higher than that of the other waveforms; the separation of the ramp waveform function from the remaining functions was less marked for this subject than for the other two subjects, however. These differences between subjects are supported by the different pattern of significant differences in the relevant Transition Points (see Appendix C).

3.7.5 Discussion

Transition Points are dependent on waveform type. For corrugation frequencies of 0.4 - 1.6 cpd, Transition Points increase in the following order: sine wave, triangle wave, ramp, square wave. This ordering highlights the necessity of specifying the type of motion to be judged: Ono and Steinbach (1990) found that the relative amounts of motion perceived in these surfaces was in the opposite direction to that observed in the current experiment, but noted that the type of motion was not identical across waveforms.

The relative ordering of the waveforms' Transition Points precludes a characterisation of the processes underlying the Transition Point as a 'channels-based' one. If this were the case then the Transition Points of the sine wave would be greater than those of the other compound stimuli with the same fundamental frequency. As an illustration, consider the amplitudes of the components of the square wave of corrugation frequency 0.05 cpd at its Transition Point, 16.146 arc min equivalent disparity, for Subject SAS. Given that

$$sq(x) = 4/\pi \{(2\pi f) + 1/3\ (2\pi^*3f) + 1/5\ (2\pi^*5f) + .. + 1/n\ (2\pi^*nf)\},$$

the amplitudes (arc min equivalent disparity) of the sinusoidal

components are

0.05 cpd : $4/\pi$ (16.416) = 20.902
0.15 cpd : $4/\pi$ (16.416*1/3) = 6.967
0.25 cpd : $4/\pi$ (16.416*1/5) = 4.180
0.35 cpd : $4/\pi$ (16.416*1/7) = 2.986
0.45 cpd : $4/\pi$ (16.416*1/9) = 2.322

.
.

Apart from the 0.05 cpd surface, the Transition Points of sine waveforms of the above corrugation frequencies were not measured. The Transition Points (arc min equivalent disparity) of measured corrugation frequencies are shown below.

0.05 cpd = 8.804
0.1 cpd = 4.613
0.2 cpd = 3.289
0.4 cpd = 1.143.

On the basis of these results it is expected that the Transition Points of sine waveforms of 0.15, 0.25, 0.35 and 0.45 cpd would be considerably *lower* than the amplitudes of these component frequencies of the square waveform at its Transition Point. The amplitude of the 0.05 cpd component (20.902 arc min equivalent disparity) is substantially greater than the Transition Point of the surface containing this frequency alone (8.804 arc min equivalent disparity).

Experiments 1a, 1b and 2, and the work of Welch (1996), Sachtler and Zaidi (1995) suggest that the motion gradients or the curvature (Norman et al., 1991) of a surface may be more important than the Fourier components in determining how the surface is perceived. The similarity of the sine and triangle wave Transition Points (at least at the lower frequencies) supports this notion. These waveforms are similar in appearance and contain similar MMGs, although that of the sine waveform is slightly greater than that of the triangle waveform. This is reflected in the slightly lower Transition Points of the former waveform.

The MMGs of the square and ramp waveforms can not account for the Transition Points, however, because for these surfaces the MMGs are infinitely steep and thus extremely low Transition Points would be predicted. The results are, however, in qualitative agreement with an analysis which takes into account both the maximal and the non-maximal motion gradients. Large regions of the square wave surface contain a motion gradient of zero, which would be expected to result in infinitely high Transition Points. The presence of the extremely high gradients at the surface discontinuities would be expected to result in very low Transition Points. The observed Transition Point might be a weighted combination of the individual predicted Transition Points. This hypothesis would account for the decrease in Transition Point with increasing corrugation frequency: as frequency increases, the number of regions containing very high motion gradients also increases.

The hypothesis would also account for the relative position of the ramp waveform Transition Points. The (non-infinite) motion gradients of a ramp waveform are the same as those of a triangle wave of half the corrugation frequency and thus it would be expected that Transition Points of the ramp waveform would lie above those of the sine and triangle waveforms, a prediction borne out by the data.

In conclusion, whilst an analysis of the MMG accounts for the data of the triangle and sine waveforms, an analysis of both the maximal and non-maximal gradients provides a better account of the Transition Points of the square and ramp waveforms.

3.8 Experiment 3: Transition Point as a function of Stimulus Size

3.8.1 Introduction

Several studies have demonstrated that perspective information strongly influences the perception of rotation of parallax stimuli (e.g. Gibson and Gibson, 1957; Braunstein and Payne, 1968; Börjesson, 1971; Hershberger, Stewart and Laughlin, 1976; Braunstein, 1977a; 1977b; Rogers and Rogers, 1992), and of slant in stationary, stereoscopic surfaces (e.g.

Epstein and Mountford, 1963; Clark, Smith and Rabe, 1966; Youngs, 1976; Stevens and Brookes, 1988; Gillam and Ryan, 1992; Eby and Braunstein, 1993; Ryan and Gillam, 1994). Specifically, studies have shown that perspective has a strong role in quantitative estimates of slant angle (Gibson and Gibson, 1957) and perceived rotation (Williams, 1993). These studies are considered in Chapter 6, where a discussion of the geometry of perspective information is additionally provided.

Petersik (1991) asked observers to estimate the extent of perceived rotation of a sphere under perspective projection. The sphere radius was either 4.45 or 8.9 cm, and the simulated rotation magnitude was between 25° and 65°. The magnitude of estimated rotation was consistently underestimated, but it was found that for a given angle of rotation, this underestimation was greater for the smaller than for the larger sphere. This effect was not attributable to the greater linear velocities contained in the larger stimulus. A possible reason for the larger sphere's greater apparent rotation is that the absolute size of the perspective cue was greater in this stimulus than in the smaller stimulus and thus affected perceived rotation to a greater extent.

The present experiment examined the effect of size on the perception of threshold levels of rotation. The absence of vertical perspective is a strong cue signalling a *lack* of rotation. Any percept of rotation would obviously involve a conflict with the information provided by vertical perspective (given that the stimulus is large enough for the perspective information to be detectable). The absolute magnitude of perspective change accompanying a given rotation angle of a surface viewed from a given distance increases with increasing stimulus size. In the absence of perspective change therefore, the perception of rotation in large stimuli involves a greater conflict than the perception of rotation in small stimuli, because the difference between the expected and the observed amounts of perspective change is greater in the former case than in the latter. On this basis it is predicted that rotation will be more easily perceived as stimulus size decreases. Thus it is expected that Transition Points will decrease with decreasing stimulus size.

3.8.2 Stimuli

Stimuli were of sizes 5, 10 and 20 °. The corrugation frequency was 0.2 cpd.

3.8.3 Results

Figs 3.23 - 3.25 show Transition Point as a function of stimulus size. Lines of power least squares fit ($y = ax^b$), providing linear fits on log-log axes, are fitted to the data. It can be seen that there is a slight increase in Transition Points with increasing stimulus size. The slope of this function is similar for all three subjects. Each doubling of size results in an increase in the Transition Point by a factor of approximately 0.2.

3.8.4 Discussion

Transition Points increase with increasing stimulus size. If stimulus size was directly related to the Transition Point then it would be expected that the function would have a gradient of +1. This is clearly not the case. The gradient is of the order ~ +0.25, demonstrating that the Transition Point is biased by, but not determined by, the stimulus size.

This result is consistent with the notion that the effects of vertical perspective on rotation perception exist at threshold levels of rotation, suppressing the perception of rotation to an increasing extent as stimulus size is increased. The ability of linear perspective cues specifying zero slant to successfully compete with additional cues specifying non-zero slant has also been demonstrated using stereoscopic surfaces. Gillam and Ryan (1992) asked subjects to adjust the slant of a disc so that it matched the slant of a surface which was stereoscopically rotated about either the horizontal or the vertical axis. The surface markings consisted of either horizontal lines, vertical lines, lines of both orientations, a diagonal grid, a grid containing horizontal, vertical and diagonal lines, or random dots. The surface markings were consistent with a frontal surface. It was found that slant about the vertical axis was underestimated to the greatest extent when the surface markings contained horizontal lines, demonstrating that the absence of linear perspective can compete especially successfully with conflicting stereo cues which specify that the surface is slanted. A follow-

up study (Ryan and Gillam, 1994) showed that as horizontal lines were successively added to a surface which was stereoscopically rotated about its vertical axis and which initially contained only vertical lines, the magnitude of perceived slant decreased with the addition of each line. The study also found that the implicit contours caused by regular line endings (such as the implicit horizontal contours created by regularly spaced vertical lines of equal height) are also effective inhibitors of perceived slant about the vertical axis.

Stevens and Brookes (1988) showed observers surfaces in which the direction of slant specified by the stereo disparity gradient was in the opposite direction to that specified by the monocular perspective cues. Two dots, the probe and the reference, were present on the surface. Subjects judged whether the probe dot was nearer or further away than the reference dot. It was found that 92% of the relative depth judgements were consistent with the perspective information. Cagenello and Rogers (1993) investigated the influence of surface markings on stereoscopic slant thresholds. They found that thresholds were higher (by a factor of 2.5 to 3.5) when the surface was marked with lines of 0°/90° than when it was marked with lines of ±45°. This difference was attributed to the differences in orientation disparity generated by these two sets of markings. Although deemed unlikely by the authors, it is possible that the relative strengths of the conflicting perspective information provided by these two sets of lines also contributed to the results. Lines of 0°/90° provide a stronger cue that a surface is frontal than lines of ±45° (e.g. Gillam and Ryan, 1992).

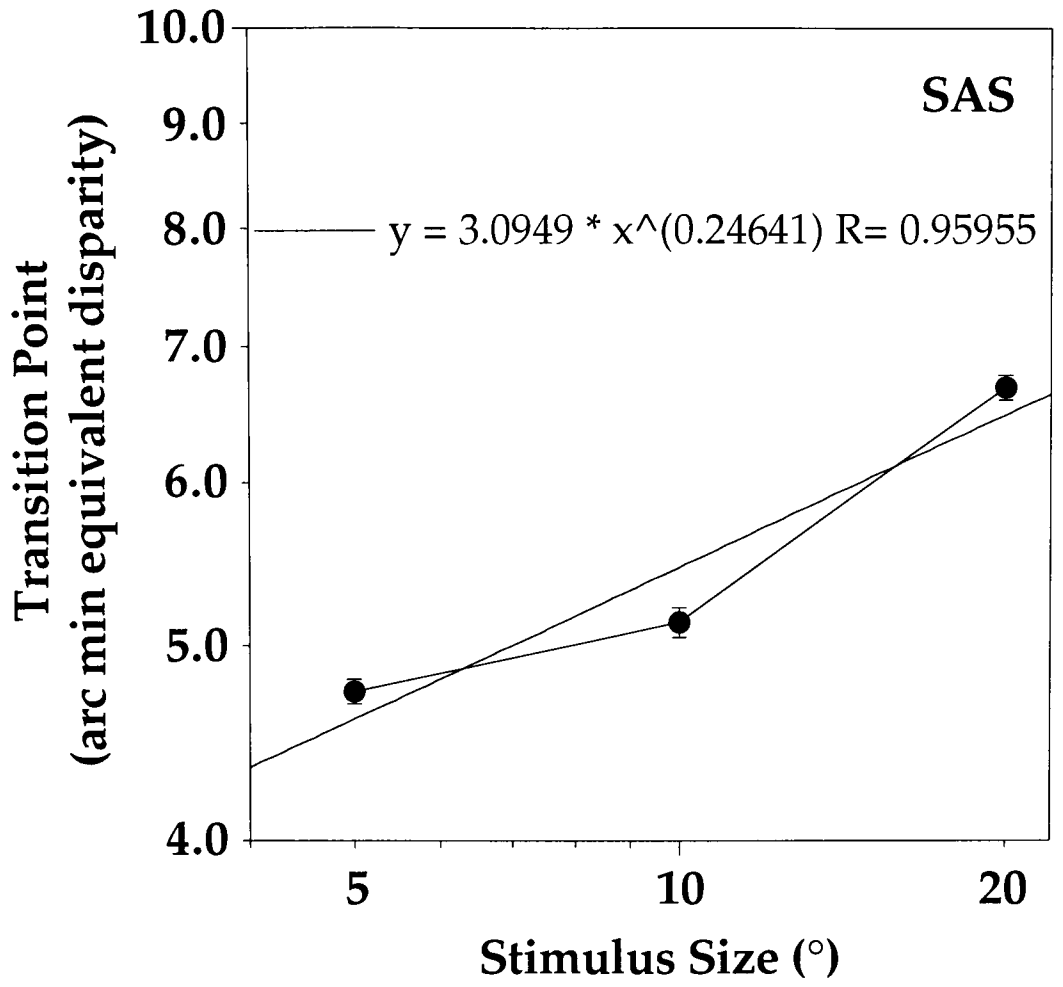


Fig. 3.23: Transition Point as a function of Stimulus Size for Subject SAS.

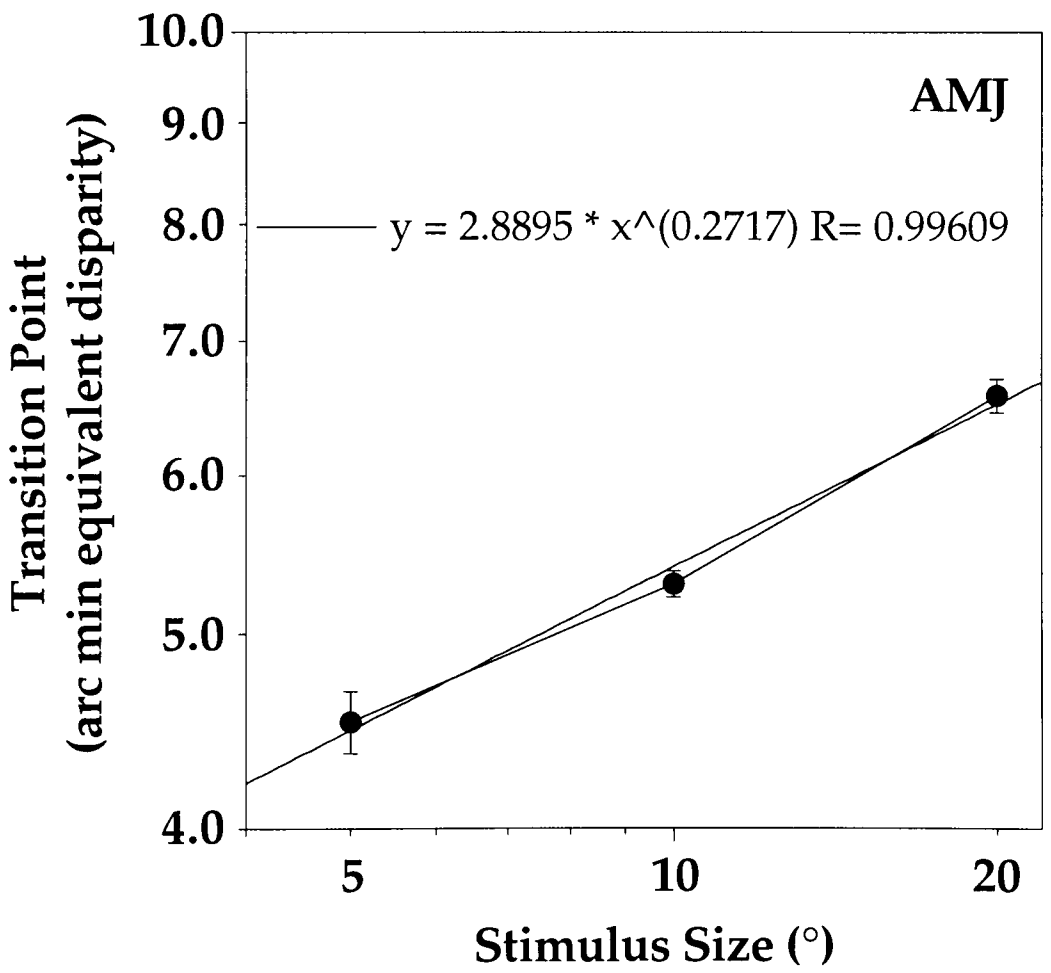


Fig. 3.24: Transition Point as a function of Stimulus Size for Subject AMJ.

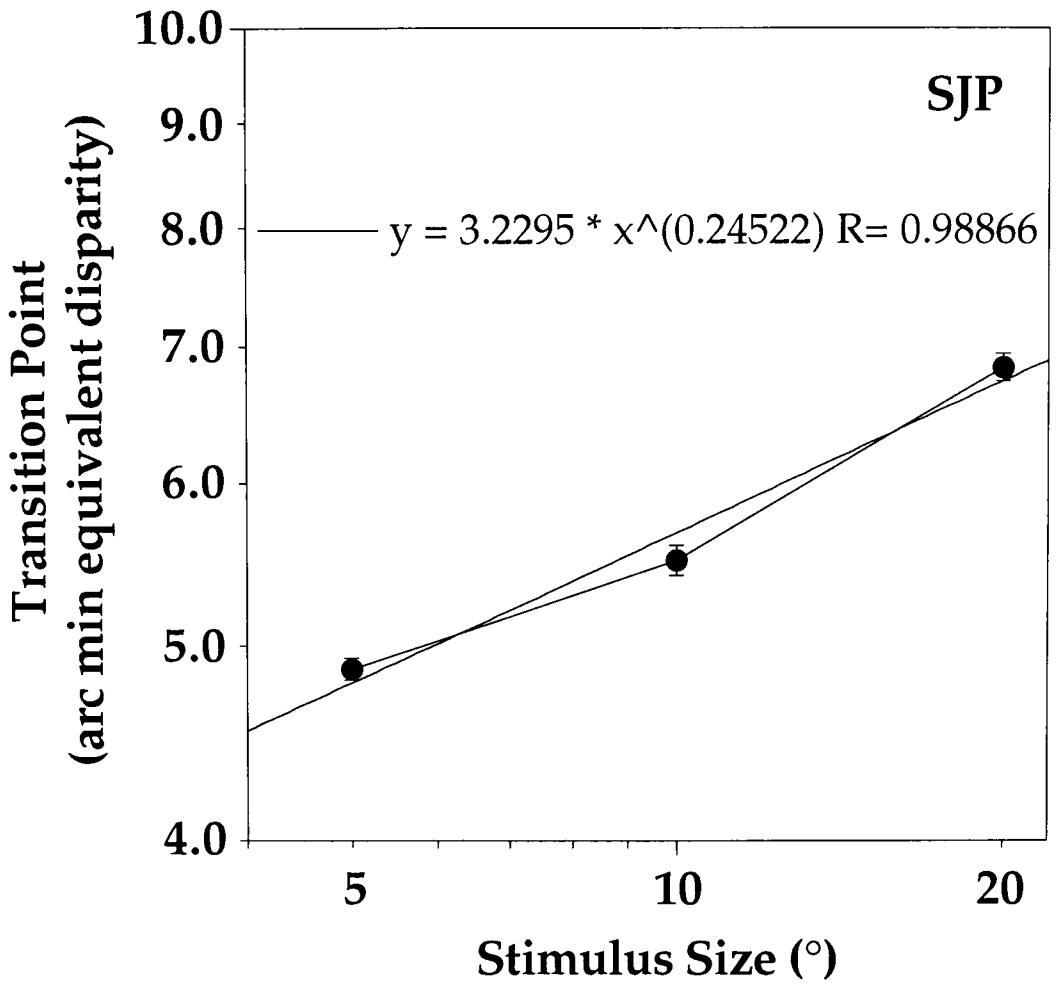


Fig. 3.25: Transition Point as a function of Stimulus Size for Subject SJP.

The influence of surface markings on stereoscopic slant thresholds was also investigated by Mitchison and McKee (1990). For five of the six observers, thresholds were higher by a factor of 1.88 when the markings were 0°/90° than when they were ±45°. The relative sizes of the differences between the thresholds of the two sets of markings in this study and that of Cagenello and Rogers (1993) may be partly attributable to the different stimuli sizes used. The latter study used stimuli which were 10.66° in diameter whereas stimulus size in the present study was 0.75°. McKee (1983) found that the addition of horizontals to a pair of vertical lines, thus forming a square, greatly increased the stereo thresholds for the relative depths of the verticals. Arditi (1982) found that thresholds for slant detection were lower for a ±45° cross hair pattern than for a 0°/90° cross hair pattern.

In sum, these studies illustrate that conflicting perspective information can affect the slant perception of stereoscopic surfaces even when slant is present at threshold levels. The results of the current experiment demonstrate that the presence of perspective information specifying that the surface is in the frontal plane can bias the rotation threshold of motion

parallax surfaces.

3.9 Experiment 4: Transition Point as a function of Viewing Distance

3.9.1 Introduction

Ono et al. (1986) found that the frequency with which observers perceived parallax surfaces to rotate increased with increasing viewing distance (see Section 3.2.3). However, as noted above, this study did not use stimuli which were proximally identical. As viewing distance increased, stimulus size and corrugation frequency remained distally constant, such that their proximal values were respectively halved and doubled with each doubling of viewing distance. Experiments 1a and 3 have demonstrated that each of these factors affects the Transition Point in a direction which would predict the results of Ono et al. (1986). The current experiment examines the Transition Points of proximally identical stimuli, in order that the confounding variables are eliminated.

3.9.2 Stimuli

Stimuli were proximally identical in terms of their size, corrugation frequency and element size. Stimuli were of size 5°. The corrugation frequency was either 0.2 or 0.4 cpd.

3.9.3 Procedure

There were three viewing distances: 57, 114 and 228 cm. The order of the experimental blocks (see Section 3.4) was counter-balanced over the three viewing distances. The procedure was otherwise identical to that of previous experiments.

3.9.4 Results

Figs. 3.26 and 3.27 show Transition Points as a function of viewing distance for each corrugation frequency. The open squares and closed circles represent Transition Points for the 0.2 and 0.4 cpd surfaces

respectively. For each corrugation frequency, Transition Points decreased with increasing viewing distance. Power functions of the form $y = ax^b$ have been fitted to the data. The parameters of these functions are shown on each graph.

3.9.5 Discussion

Transition Points decrease with increasing viewing distance. The slopes of the functions lie between -0.51 and -0.72, indicating that Transition Points are not solely dependent upon viewing distance. This finding is in qualitative agreement with Ono et al.'s (1986) report that the frequency of perceived rotation increased with increasing viewing distance. There is no obvious reason for this to be the case, although one hypothesis which would account for the results is that the effect of the vertical perspective information (which signalled that the surface remained in the frontal plane) on perceived rotation decreased with increasing viewing distance. In the real world, the absolute size of the vertical perspective changes which occur in a real object which rotates through a constant angle decreases with increasing viewing distance. The visual system may therefore become decreasingly reliant upon vertical perspective as a rotation cue as viewing distance increases. In the present experiment, despite the fact that the size of the vertical perspective cue remained *constant* as viewing distance increased (the stimuli were proximally identical), the putative propensity of the visual system to rely less heavily on this information because of its decreasing effectiveness (with increasing viewing distance) in the real world, may account for the observed pattern of results.

This possibility is discussed in more detail in Experiment 15, which examines the effect of viewing distance on magnitudes of perceived depth and rotation, and in Experiment 21, which examines the effect of vertical perspective on perceived rotation as a function of viewing distance.

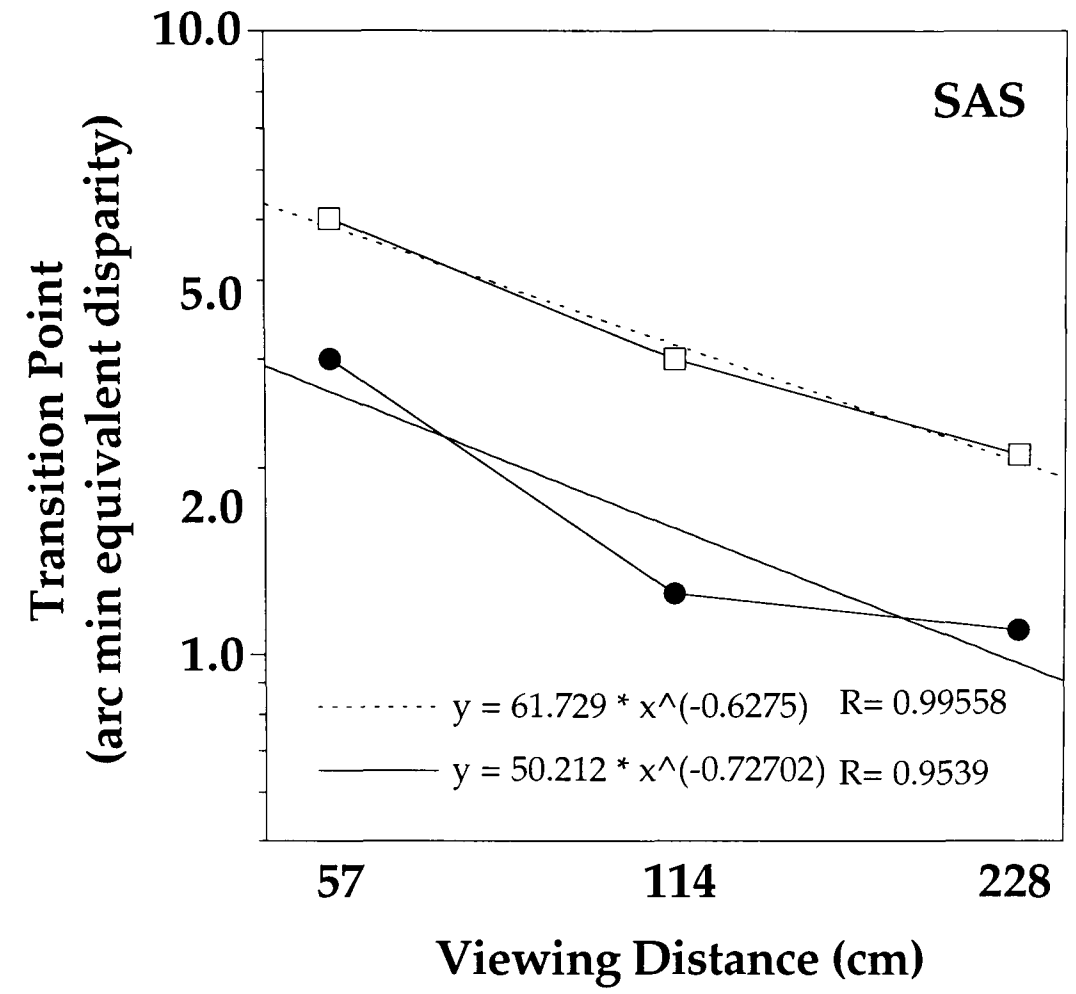


Fig. 3.26: Transition Point as a function of Viewing Distance for subject SAS. Open squares and closed circles represent data for corrugation frequencies of 0.2 and 0.4 cpd respectively.

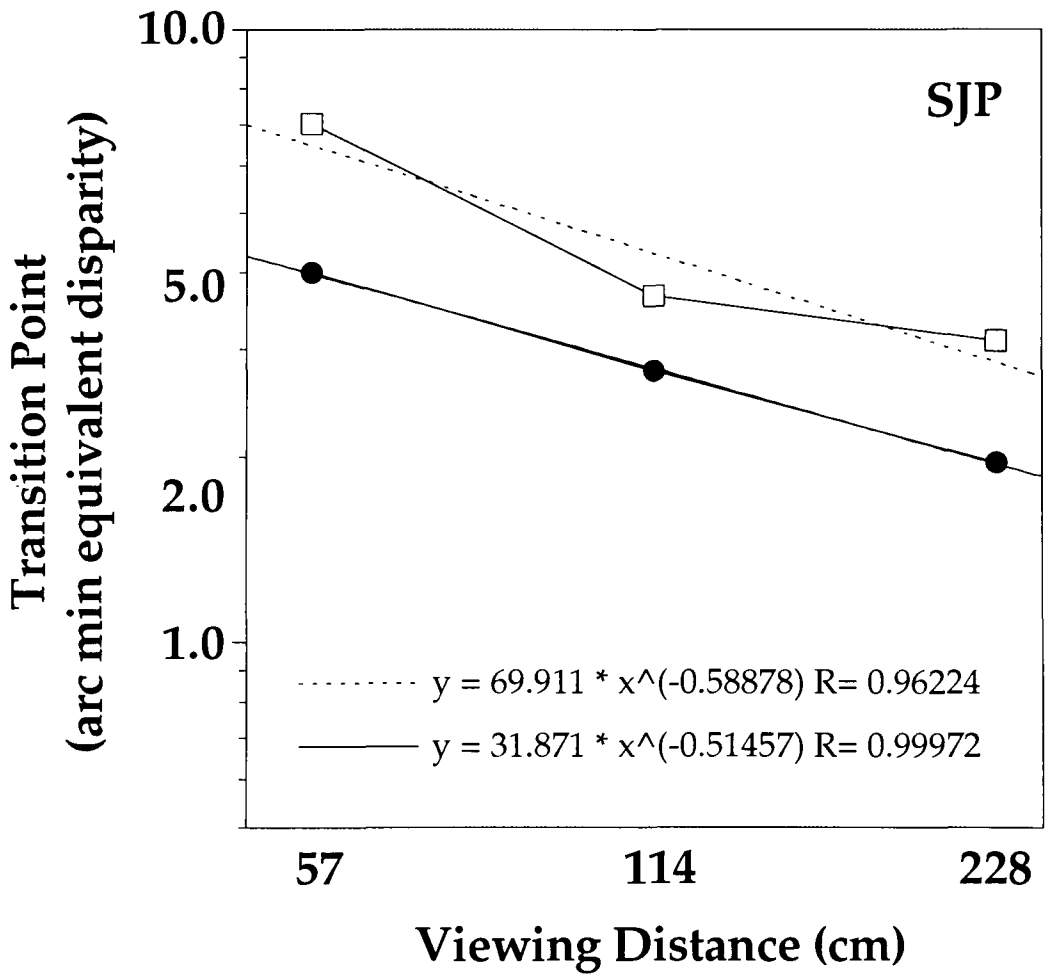


Fig. 3.27: Transition Point as a function of Viewing Distance for subject SJP. Open squares and closed circles represent data for corrugation frequencies of 0.2 and 0.4 cpd respectively.

3.10 Experiment 5: Perceptual Zones into which the percept arising from a motion parallax stimulus can fall

3.10.1 Introduction

Section 3.1 discussed the qualitatively different percepts which can arise from a motion parallax stimulus, and introduced the concept of perceptual zones separated by perceptual thresholds. Three such zones were mentioned: the zone in which relative motion amplitude is beneath the depth threshold, resulting in the percept of a stationary, 2-D surface, the zone in which a rigid, stationary, depth-modulated surface is perceived, and finally the zone in which the rigid (sinusoidal) surface appears to be rotating. These are represented in Fig. 3.28. The dashed lines dividing the zones represent the perceptual thresholds.

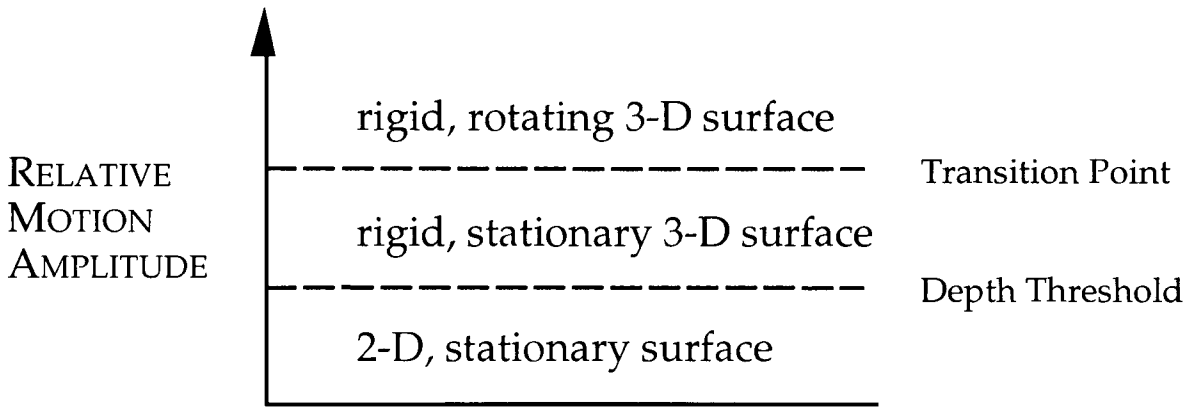


Fig. 3.28: Representation of the perceptual zones discussed in Section 3.1.

Saida and Ono (1989) provided evidence suggestive of the existence of a further zone. Observers were asked to indicate the amount of perceived depth and motion in parallax stimuli containing one of five different relative motion values. At amplitudes greater than 60 arc min equivalent disparity, depth order became unstable and the amount of perceived depth decreased. Ono et al. (1986) also found a breakdown in the 3-D percept: at a viewing distance of 320 cm observers perceived only 2-D shearing motion in a parallax display. These studies indicate the presence of a zone in which the parallax stimulus no longer gives rise to a stable, 3-D percept.

Ono et al. (1986) reported that observers perceived depth and 'rocking' motion at a viewing distance of 160 cm, but perceived 2-D shearing motion at 320 cm. It is therefore probable that there are a range of percepts incorporating both a 'rocking' motion *and* 2-D relative motion. The latter

motion would be incorporated into the 3-D percept as a relative shearing between the peaks and troughs. This surface would obviously be perceived as slightly non-rigid.

Fig. 3.29 includes the further perceptual zones suggested by the above studies. The perceptual border dividing a rigid, rotating 3-D surface and a non-rigid 3-D surface is termed the Upper Transition Point, and the border dividing a non-rigid 3-D surface and 2-D shearing motion is termed the 2-D Shearing Threshold.

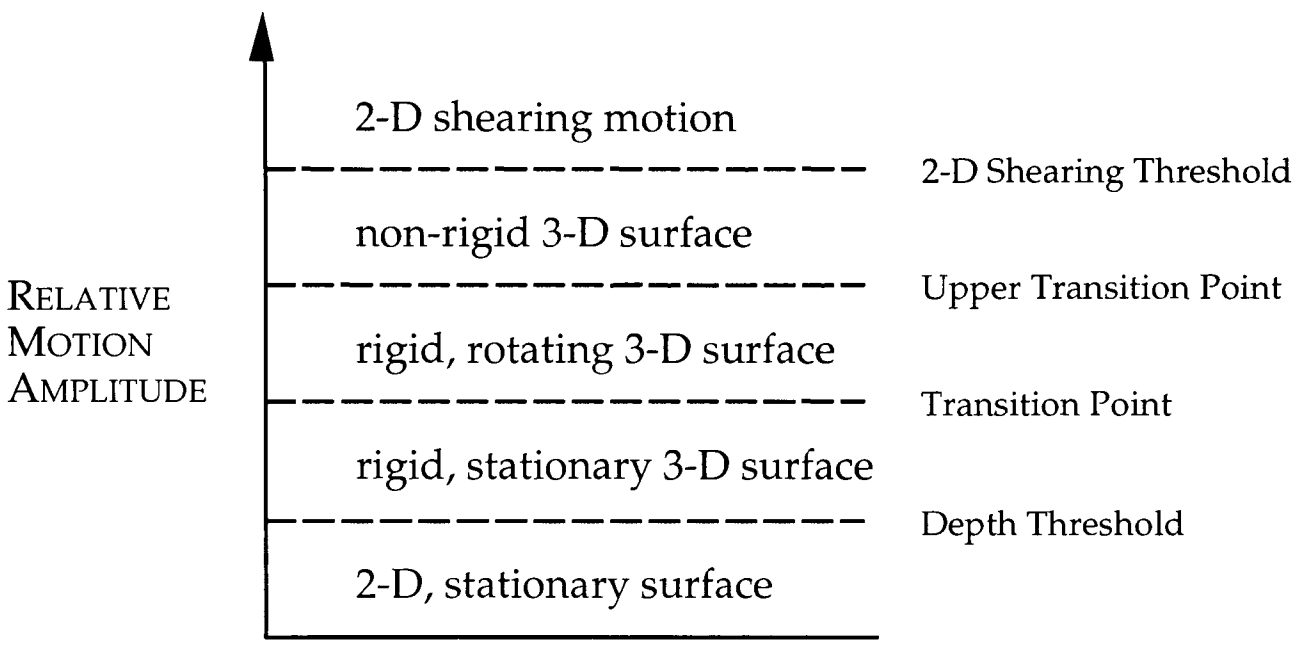


Fig. 3.29: Representation of the complete range of perceptual zones arising from a sinusoidally corrugated motion parallax stimulus.

Although numerous studies have commented that various factors influence the nature of the percept arising from motion parallax stimuli, work has generally focused upon a single border. No work to date has specifically addressed, in a single experiment, the various perceptual zones into which a parallax stimulus can fall. The aim of the present experiment was to map the various perceptual zones arising as a result of variations in the relative motion amplitude, as a function of corrugation frequency. To this end, Depth Thresholds, Upper Transition Points and 2-D Shearing Thresholds were determined in order that these data could be combined with the previously determined Transition Points.

3.10.2 Stimuli

Stimuli were sinusoidally depth modulated surfaces of corrugation frequencies ranging from 0.05 - 1.6 cpd. Stimulus size was 20°.

3.10.3 Procedure

To determine the Upper Transition Point observers were asked to decide whether the surface appeared to be either rigid and rotating, or non-rigid. The method of adjustment was used. The subject used a potentiometer to adjust the amplitude of relative motion until the percept of a rigid and rotating surface changed to one of a non-rigid surface, or vice versa. Subjects were encouraged to make initially coarse adjustments in order to generously bracket the Upper Transition Point, and to make increasingly fine adjustments until they were satisfied that the Upper Transition Point had been located. At this point the subject pressed a button, which resulted in the recording of the relative motion amplitude. Six estimations of each Upper Transition Point were made. Again, as noted in Section 3.4 with regard to the Transition Point, this estimation was necessarily subjective: there *is* no objective measure of this point. Again, however, subjects were instructed to try to maintain the same perceptual criteria across conditions.

The same procedure was employed to determine the 2-D Shearing Threshold. However, see Section 3.10.4 below.

In a further set of trials Depth Thresholds were determined. The method of constant stimuli was used. The experimental session for each corrugation frequency consisted of 280 trials in four blocks. Each block contained 70 trials, corresponding to a total of 40 trials at each of the stimulus levels. The order of blocks was randomised.

3.10.4 Results

Figs 3.30 and 3.31 show Depth Thresholds, Transition Points, Upper Transition Points and 2-D Shearing Thresholds for two subjects.

The Upper Transition Point fell with increasing corrugation frequency. The slope of the function was displaced above, and was slightly shallower than, the function of the Transition Point between a rigid, stationary surface and a rigid, rotating surface, but was otherwise similar. The 2-D Shearing Thresholds lay above the maximum relative motion amplitude possible with the apparatus (80 arc min equivalent disparity).

Depth thresholds had the characteristic U-shape (Rogers and Graham, 1982). The areas between the functions represent perceptual zones – ranges of relative motion amplitudes that support the given percept. Stimuli with amplitudes falling in the dotted area give rise to a percept of a rigid, stationary structure; the hatched area represents the amplitude range that results in the percept of a rigid, rotating structure. Relative motion values greater than the Upper Transition Point give rise to non-rigid structures. The two shaded areas together represent the range of amplitudes that are interpreted as rigid, 3-D structures.

3.10.5 Discussion

Two main observations may be made. Firstly, it is evident that at low corrugation frequencies, stimuli are interpreted as rigid, 3-D structures over a relatively large range of relative motion amplitudes (e.g. at 0.05 cpd, rigidity is perceived between amplitudes of 1 - 2 and 11 arc min equivalent disparity). At higher corrugation frequencies the range is much reduced: at 1.6 cpd it is just 3 - 5 arc min equivalent disparity. The reason for this reduced range at high corrugation frequencies is the 'U' shape of the depth threshold function in conjunction with the decrease in the Upper Transition Point.

Secondly, it is notable that at the highest corrugation frequency examined (1.6 cpd), there is no zone representing the rigid, stationary percept: as soon as there is enough relative motion to support the perception of depth, (i.e. as soon as the depth threshold is reached) the observer typically perceives the structure as rotating. Williams (1993) investigated the threshold at which observers perceived non-rigidity in monocular motion parallax stimuli of corrugation frequency 0.28 cpd, and obtained a value of approximately 12 arc min equivalent disparity. This result is similar to those obtained in the present experiment. Caudek and Proffitt (1993) presented observers with a simulated dihedral angle — two slanted planes meeting at a central horizontal ridge. There were two velocity flow field conditions. In one condition the flow field contained both object-relative and observer-relative information; in the other condition the flow field contained only object-relative information. Observers were asked to rate the apparent rigidity of the structure on a scale between 0 and 1. The projected width and the simulated depth of the

wedges were systematically varied. It was found that perceived rigidity decreased firstly as simulated depth increased, and additionally as the wedge width decreased. Both of these results are in accord with the observed variation in the Upper Transition Point with corrugation frequency. It can be seen that the zone which encompasses the greatest range of relative motion amplitudes is that in which non-rigidity is perceived.

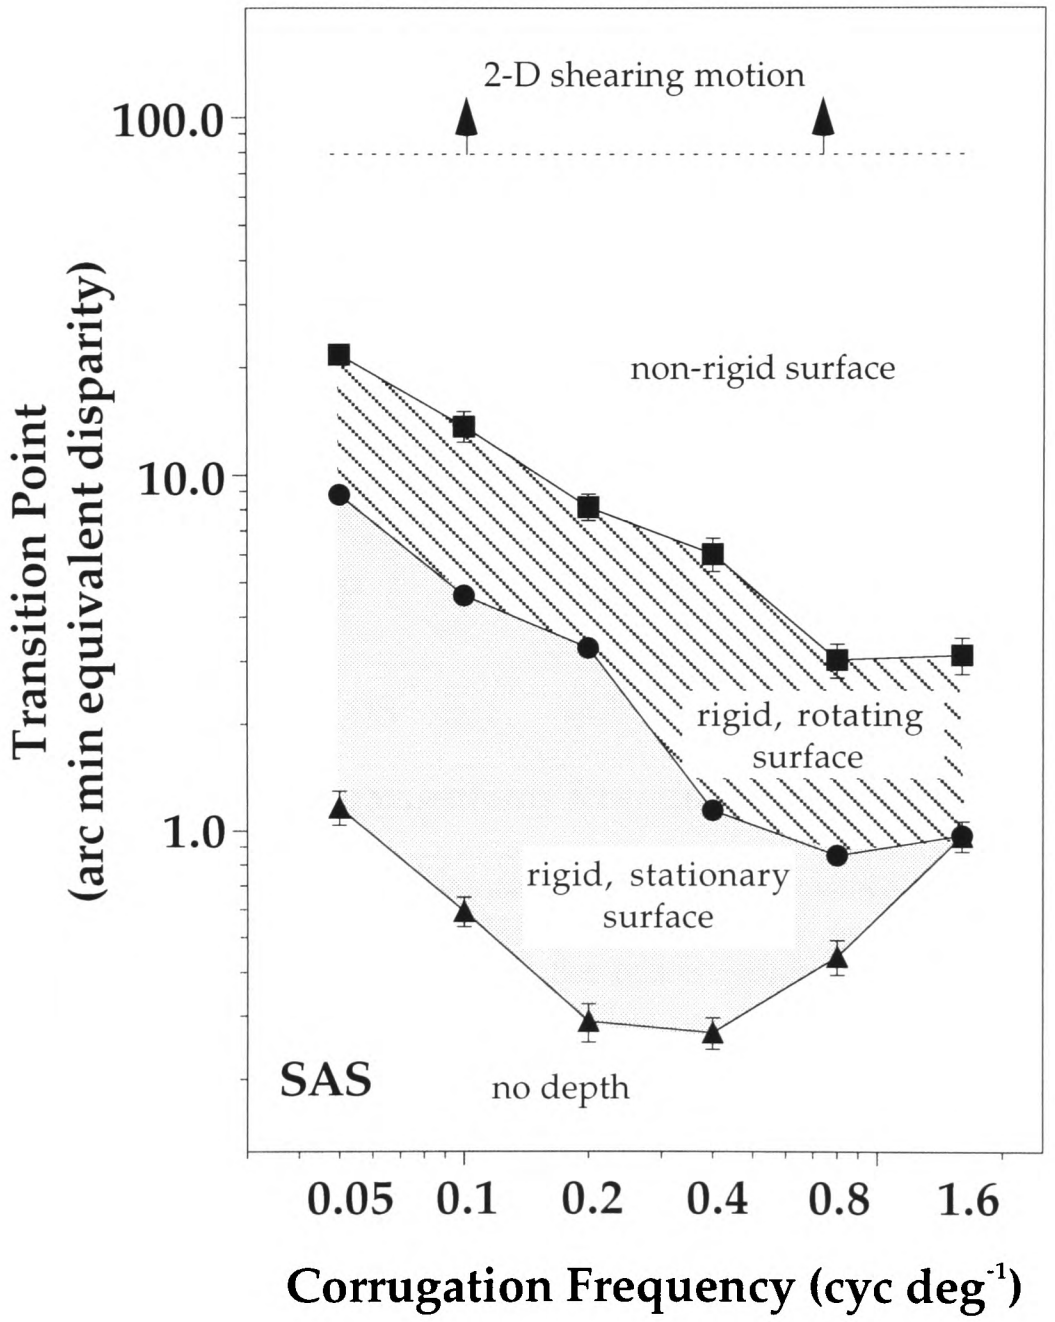


Fig. 3.30: Depth Threshold (filled triangles),Transition Point (filled circles) and Upper Transition Point (filled squares) as a function of Corrugation Frequency, for Subject SAS. See text for explanation of shaded areas.

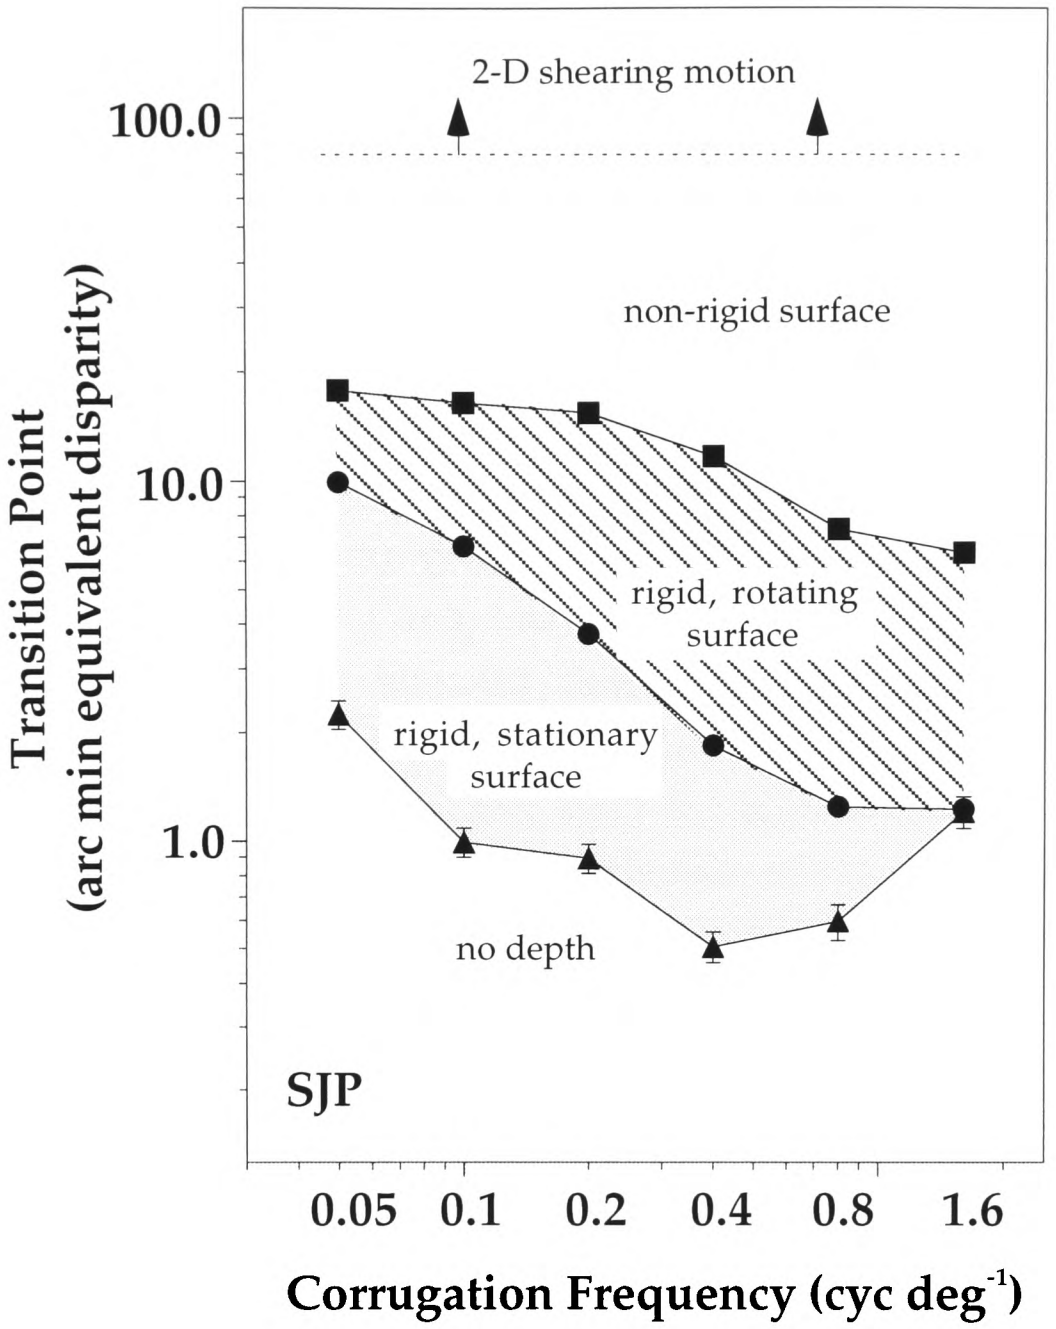


Fig. 3.31: Depth Threshold (filled triangles),Transition Point (filled circles) and Upper Transition Point (filled squares) as a function of Corrugation Frequency, for Subject SJP. See text for explanation of shaded areas.

3.11 Summary and Conclusions

The Transition Point between a rigid, stationary surface and a rigid, rotating surface is affected by a variety of factors. Experiments 1a and 1b demonstrated that an important determinant of the Transition Point is the MMG present in the surface. The former experiment showed that over a four-octave range of corrugation frequencies (three-octave for one subject) the Transition Point fell approximately linearly with increasing frequency. Experiment 1b extended this finding by demonstrating that stimuli of different corrugation frequencies but identical MMGs had very similar Transition Points. Both experiments additionally showed that beyond the motion gradient-dependent range, the Transition Point

becomes dependent upon the absolute amount of relative motion in the stimulus. Experiment 2 showed that of sine, triangle, ramp and square waveforms, sine waveforms have the lowest Transition Points and square waveforms have the highest. These results are reconcilable with those of Ono and Steinbach (1990) when the type of motion generally perceived in the waveforms is considered. In sine waveforms the motion is of a rotational or 'rocking' nature while the motion perceived in square waveforms is largely of a shearing nature. The Transition Points of sinusoidal and triangle waveforms are largely dependent on the MMG, whereas those of the ramp and square waveforms were shown to be dependent both on the maximal and non-maximal gradients of the surface. Experiment 3 showed that Transition Points increase with stimulus size, a finding likely to be largely attributable to the increasing conflict (with increasing stimulus size) between a percept of rotation and the lack of changes in vertical perspective which would be expected to accompany such a percept. Experiment 4 demonstrated that the Transition Point falls with increasing viewing distance.

Experiment 5 investigated the different perceptual zones into which parallax stimuli can fall. Previous work related to this issue has been limited in that it has either examined factors relating to only *one* of the borders between perceptual zones (e.g. Ono and Ujike, 1993; Ichikawa and Ono, 1996; Ujike and Ono, 1992) or has briefly commented upon the topic only as an indirect result of work addressing alternative issues in parallax (e.g. Ono et al., 1986; Ono and Steinbach, 1990). Experiment 5 determined the borders between five perceptual zones, for surfaces with corrugation frequencies between 0.05 and 1.6 cpd. It was found that the Transition and the Upper Transition Points fell with increasing corrugation frequency, meaning that for all frequencies there was a range of relative motion amplitudes supporting the percept of a rigid, rotating surface. The relative positions of the characteristic U-shaped depth threshold function and the Transition Point function means that the range of amplitudes supporting the percept of a rigid, stationary surface decreases significantly with increasing corrugation frequency, such that for the highest corrugation frequency tested (1.6 cpd) such a percept is non-existent. Observers did not lose the impression of depth in the parallax surfaces at relative motion amplitudes of up to 80 arc min equivalent disparity, illustrating that the range of amplitudes supporting the percept of a non-rigid surface is

significantly greater than the range which supports the percept of a rigid surface.

Chapter 4:

Perceived Rotation of Motion Parallax Surfaces in the Perceptual Zone above the Transition Point

4.1 Introduction

In the experiments of Chapter 3, subjects were instructed during each trial to decide whether the stimulus was 'stationary' or 'rotating', in order to determine the Transition Point between these two percepts. Whilst providing a useful examination of factors influencing the Transition Point, i.e. the *threshold* of perceived rotation, these data do not address the perception of supra-threshold rotation. The experiments of this present chapter examine the influence of various factors on the perceived rotation of surfaces in the perceptual zone above the Transition Point (see Fig. 3.28).

4.2 Experiment 6: The relative rotation of two stimuli of different corrugation frequencies

The aim of the following experiment was a preliminary analysis of perceived rotation at relative motion values greater than the Transition Point. The experiment compared the perceived rotation of two stimuli, one containing a different corrugation frequency from the other.

4.2.1 Stimuli

There were two types of stimuli, 'Standard' stimuli and 'Test' stimuli. Standard stimuli always contained the same amount of relative motion, 16 arc min equivalent disparity. Test stimuli contained one of seven

different values of relative motion; previous pilot experiments had determined a suitable range from which these values were taken. Standard stimuli contained one of three corrugation frequencies: 0.05, 0.1 or 0.2 cpd. Test stimuli contained one of four corrugation frequencies: 0.05, 0.1, 0.2 or 0.4 cpd. All stimuli were 20° in diameter. Each Standard stimulus was paired with each Test stimulus, giving a total of 12 Standard/Test pairs.

4.2.2 Procedure

The method of constant stimuli was used to determine the point at which the Test stimulus and the Standard stimulus appeared to be rotating to the same extent. On each trial observers were shown the Standard stimulus and the Test stimulus in random order. They were asked to decide which stimulus appeared to rotate more, and to press one of two buttons to record their choice. Each stimulus was presented for 2 seconds, and the inter-stimulus interval (ISI) was 0.5 seconds. During the ISI the relative motion amplitude was zero. Each experimental block contained 70 trials, 10 at each relative motion level. Four experimental blocks were carried out for each Standard/Test pair, making a total of 40 trials at each stimulus level. The order of the blocks was counter-balanced.

4.2.3 Results

Probit analysis (Finney, 1971) was used to obtain the best-fitting cumulative Gaussian curve for each data set. The 50% point was taken as the 'Rotation Match', the point at which the Test and Standard stimuli appeared to be rotating to equal extents.

Figs. 4.1 - 4.3 show the Rotation Match (expressed in arc min equivalent disparity) as a function of the corrugation frequency of the Test stimulus, for Subjects SAS, AMJ and SJP respectively. Each curve represents a Standard stimulus of a different corrugation frequency. Fig. 4.4 shows the mean results. The dotted line has a slope of -1.

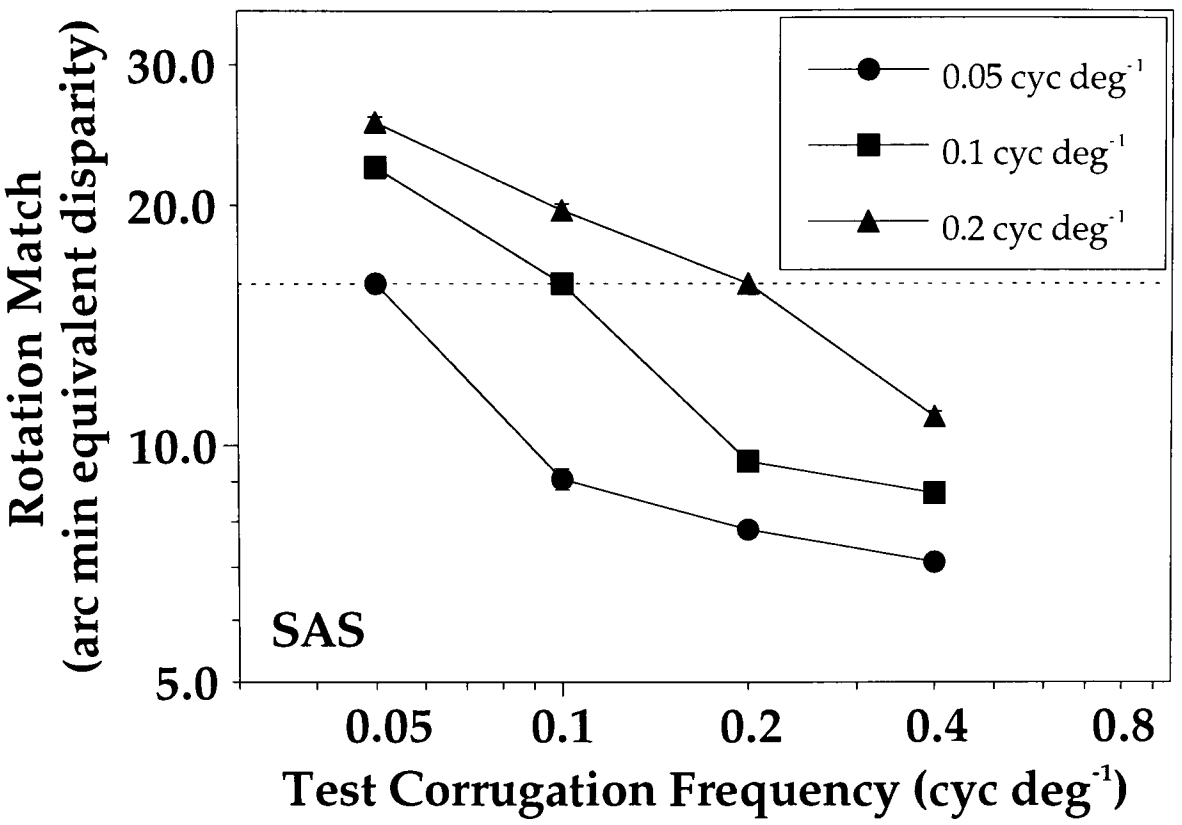


Fig. 4.1: Rotation Match as a function of Test Corrugation Frequency for Subject SAS. The three functions represent matches obtained with each of the three Standard Corrugation Frequencies, the values of which are shown in the legend. The horizontal dashed line represents the relative motion value of all Standard stimuli, 16 arc min equivalent disparity.

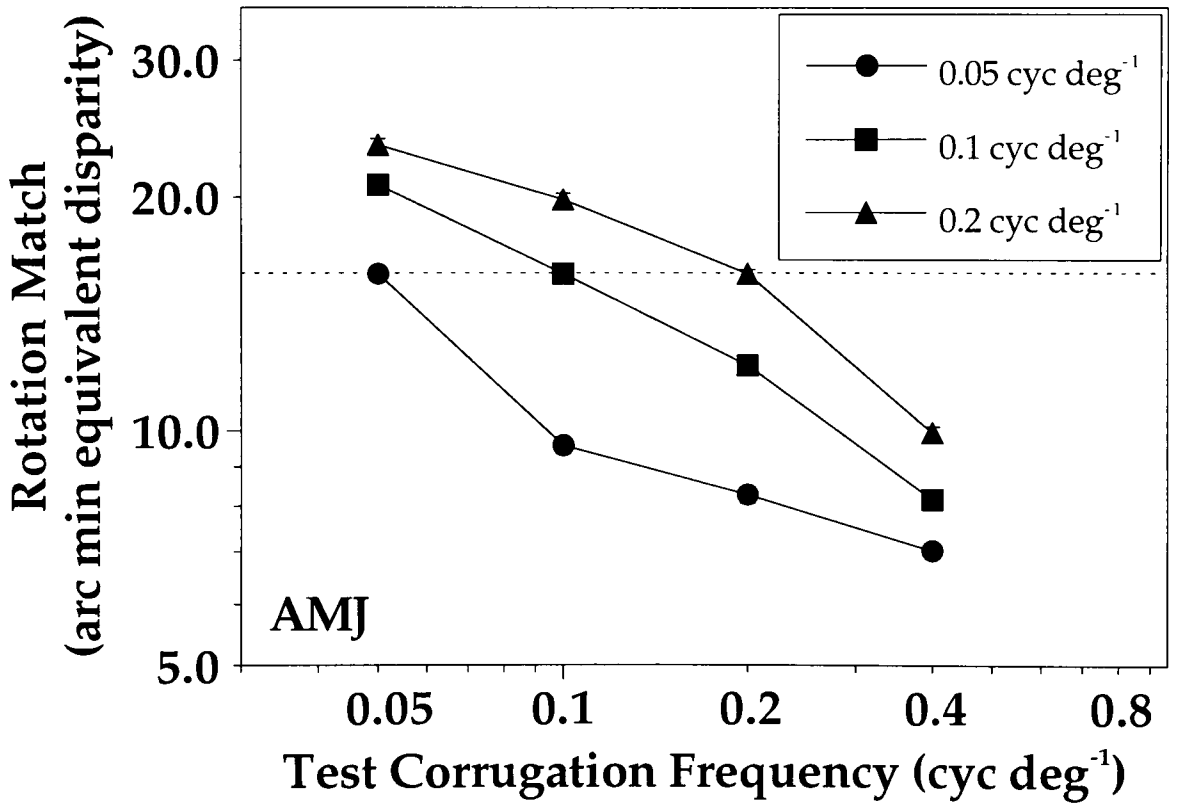


Fig. 4.2: Rotation Match as a function of Test Corrugation Frequency for Subject AMJ. The three functions represent matches obtained with each of the three Standard Corrugation Frequencies, the values of which are shown in the legend. The horizontal dashed line represents the relative motion value of all Standard stimuli, 16 arc min equivalent disparity.

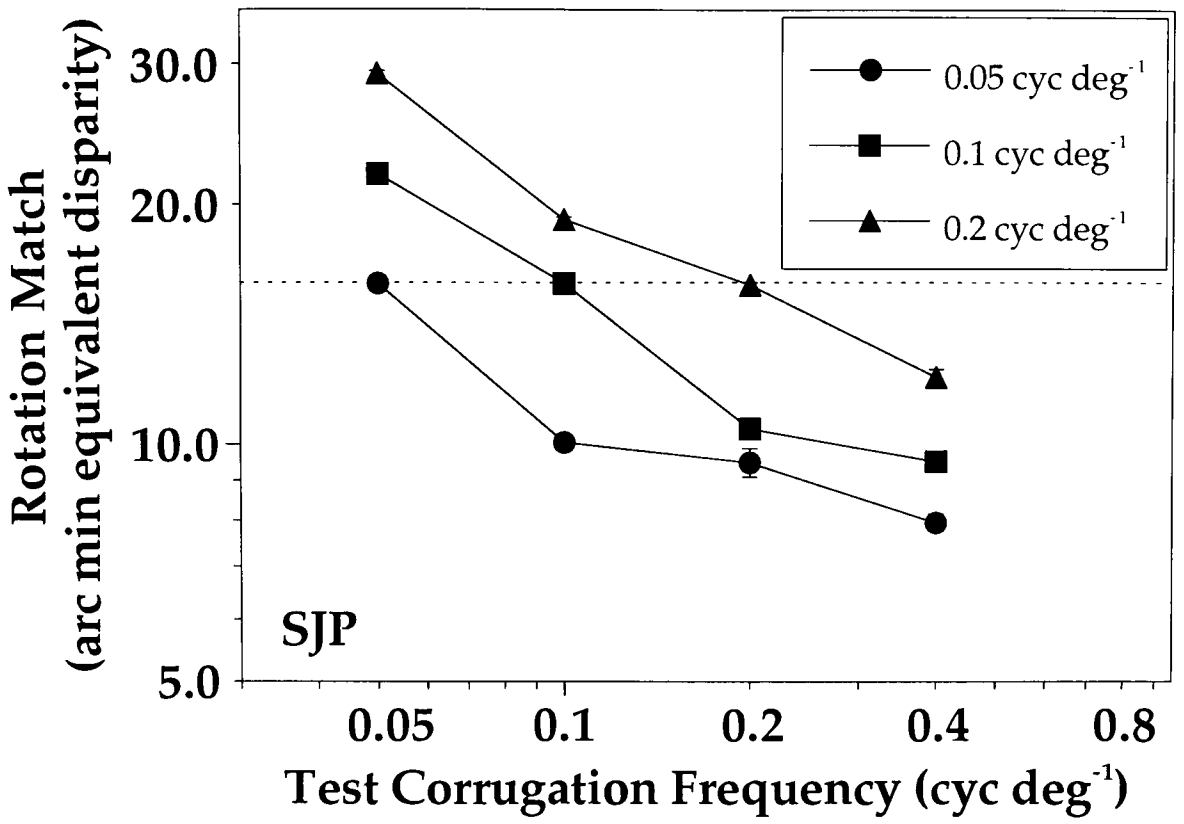


Fig. 4.3: Rotation Match as a function of Test Corrugation Frequency for Subject SJP. The three functions represent matches obtained with each of the three Standard Corrugation Frequencies, the values of which are shown in the legend. The horizontal dashed line represents the relative motion value of all Standard stimuli, 16 arc min equivalent disparity.

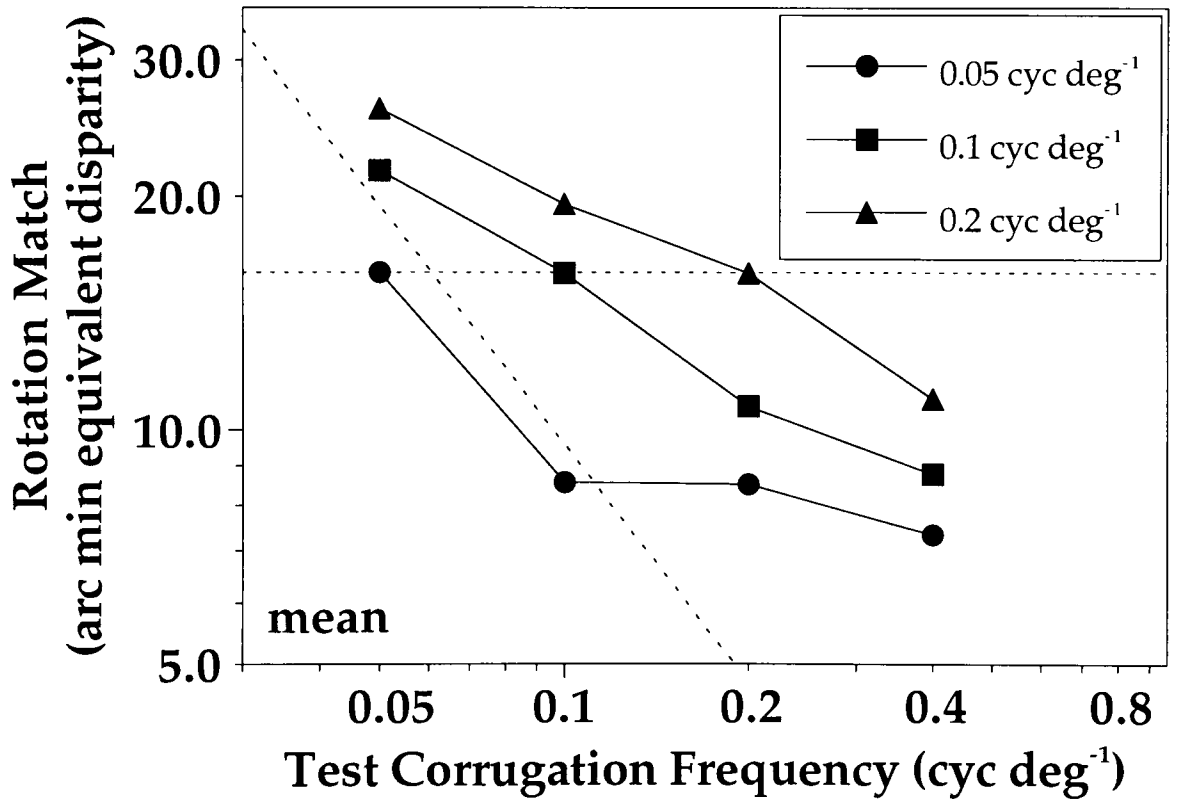


Fig. 4.4: Rotation Match as a function of Test Corrugation Frequency averaged over three subjects. The three functions represent matches obtained with each of the three Standard Corrugation Frequencies, the values of which are shown in the legend. The horizontal dashed line represents the relative motion value of all Standard stimuli, 16 arc min equivalent disparity. The sloping dashed line has a gradient of -1.

4.2.4 Discussion

Figs. 4.1 - 4.4 show that surfaces of high corrugation frequency appear to rotate more than surfaces of lower corrugation frequency. Hence, at the PSE, the relative motion amplitude of the higher frequency surface was smaller than the amplitude of the lower frequency surface. This result extends that of Experiment 1a — namely that Transition Points decrease with increasing corrugation frequency.

If perceived rotation was dependent only on the MMG of the surface, then it would be expected that the functions would have gradients of -1. It would also be predicted that the functions of the 0.05 and 0.2 cpd standard surfaces would be displaced below and above, respectively, the function of the 0.1 cpd standard surface. The gradients of all the functions were shallower than predicted, suggesting that although the MMG affects the magnitude of perceived rotation of the surface, it is not the sole determinant.

Since the relative motion amplitudes of a number of the Standard stimuli were greater than the Upper Transition Points determined in Experiment 5, it might be the case that a degree of non-rigidity was perceived in at least some of the surfaces. The experiments of this chapter are, however, concerned only with the apparent rotation of the surfaces and observers were instructed to make judgements concerning this specific attribute only. The perception of non-rigidity in the context of relative motion values which exceed the Upper Transition Point will be discussed further in Section 4.3.5.

The current experiment was useful as it provided a preliminary analysis of the relative magnitudes of perceived rotation in surfaces of various corrugation frequencies. However, the experiment was limited in that the procedure allowed only a comparison between, rather than a direct estimate of, perceived rotation magnitudes. The remainder of this chapter's experiments employed a different technique to measure perceived rotation; this is detailed in the following section.

4.3 Experiment 7a: Matched Rotation as a function of Corrugation Frequency

4.3.1 Apparatus

Rotation matches were made using a circular paddle situated at 57 cm from the subject. The bottom of the paddle was 1 cm above the top of the oscilloscope screen, and the centres of the paddle and the screen lay on the same vertical axis when the viewing distance (observer to screen) was 57 cm. The paddle was 7 cm in diameter. It was constructed of balsa wood, and its surface was covered with patterned paper. The pattern was a 50% density, irregular textured pattern of 'blobs', each of which subtended 1 - 2°. The paddle rotated about its vertical axis. The rotation of the paddle was linked to the head movement of the subject so that as the head moved anti-clockwise with respect to the screen, the paddle rotated in a clockwise direction, and as the head moved clockwise with respect to the screen, the paddle rotated in an anti-clockwise direction. The extent of rotation was controlled with a potentiometer which rested on a bench in front of the observer. Details of the calibration of the paddle are found in Section 2.1.5.

Ono and Steinbach's (1990) study, which attempted to quantify the extent of perceived motion of head-linked parallax surfaces, used a linear measure. Observers chose 'furthest' and 'nearest' parts of the stimulus and indicated the extent of perceived dot motion on a ruler. It was noted that surfaces containing motion gradients of either a ramp, triangle or sine waveform appeared to rotate. It could therefore be suggested that the 'rotation matching' procedure of the current experiment is more suited to the task at hand than Ono and Steinbach's (1990) procedure, since the motions of the matching and test stimuli are similar.

4.3.2 Stimuli

Stimuli contained one of four corrugation frequencies: 0.05, 0.1, 0.2 or 0.4 cpd. Five values of relative motion were used in each corrugation frequency condition. Stimuli were 20° in diameter.

4.3.3 Procedure

The method of adjustment was used. Subjects were asked to adjust the magnitude of rotation of the paddle so that it matched the amount of perceived rotation in the stimulus. On each trial the stimulus contained one of seven equally-spaced relative motion amplitudes. Trials were blocked in groups of fourteen, corresponding to two repetitions of each of the seven relative motion amplitudes. Subjects completed three sessions of one block per corrugation frequency per session. The order of blocks within each session was randomised.

Subjects made the rotation match using a hand-operated potentiometer, situated out of sight on a bench below eye-level. Unlimited time was allowed for each match, but observers usually required only a few seconds. In all experiments subjects were permitted to move their eyes freely around the stimulus.

4.3.4 Results

Figs. 4.5 - 4.7 show Matched Rotation as a function of relative motion for three subjects. Perceived rotation increased with increasing relative motion for all corrugation frequencies tested.

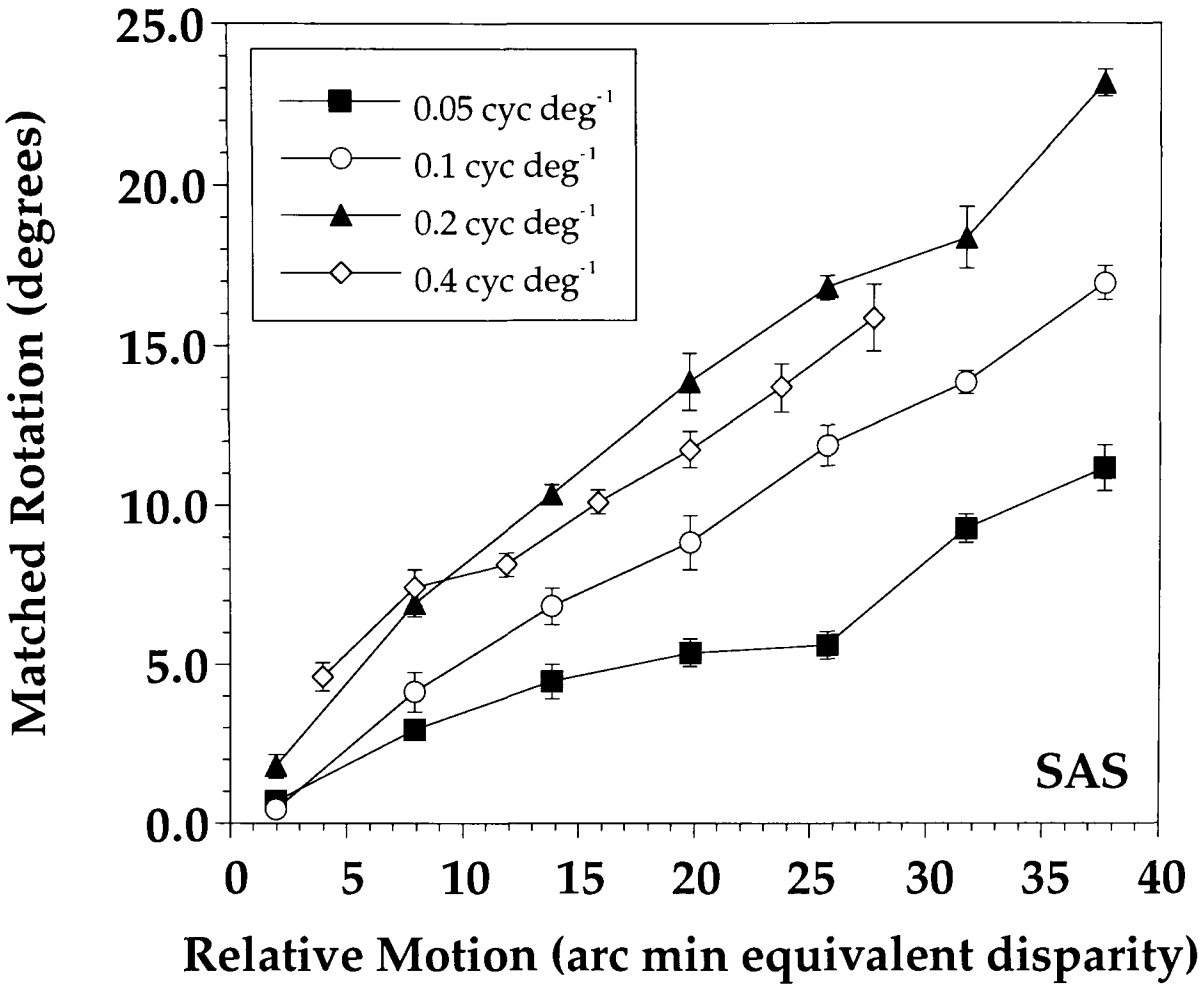


Fig. 4.5: Matched Rotation as a function of Relative Motion for four Corrugation Frequencies for Subject SAS.

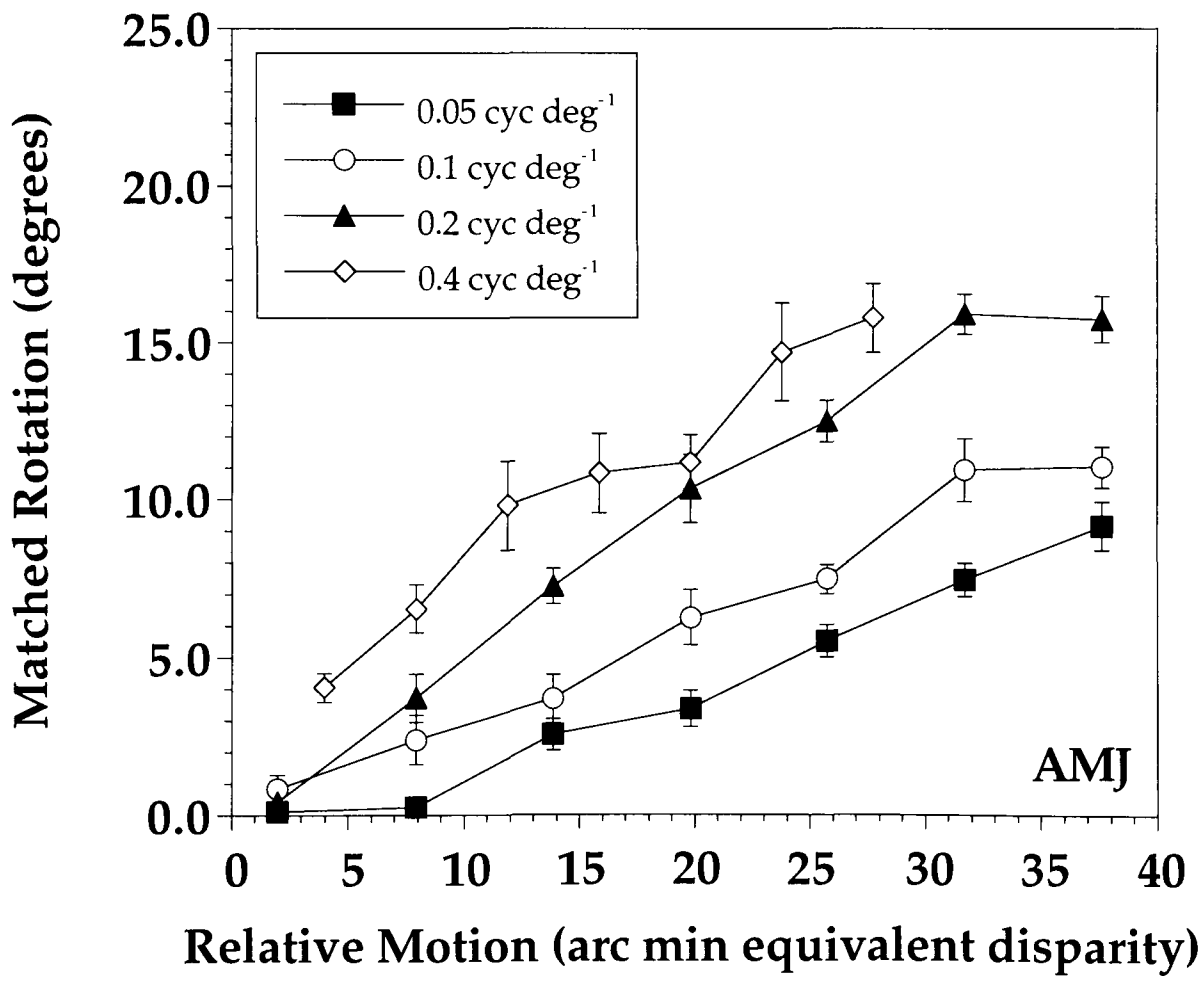


Fig. 4.6: Matched Rotation as a function of Relative Motion for four Corrugation Frequencies for Subject AMJ.

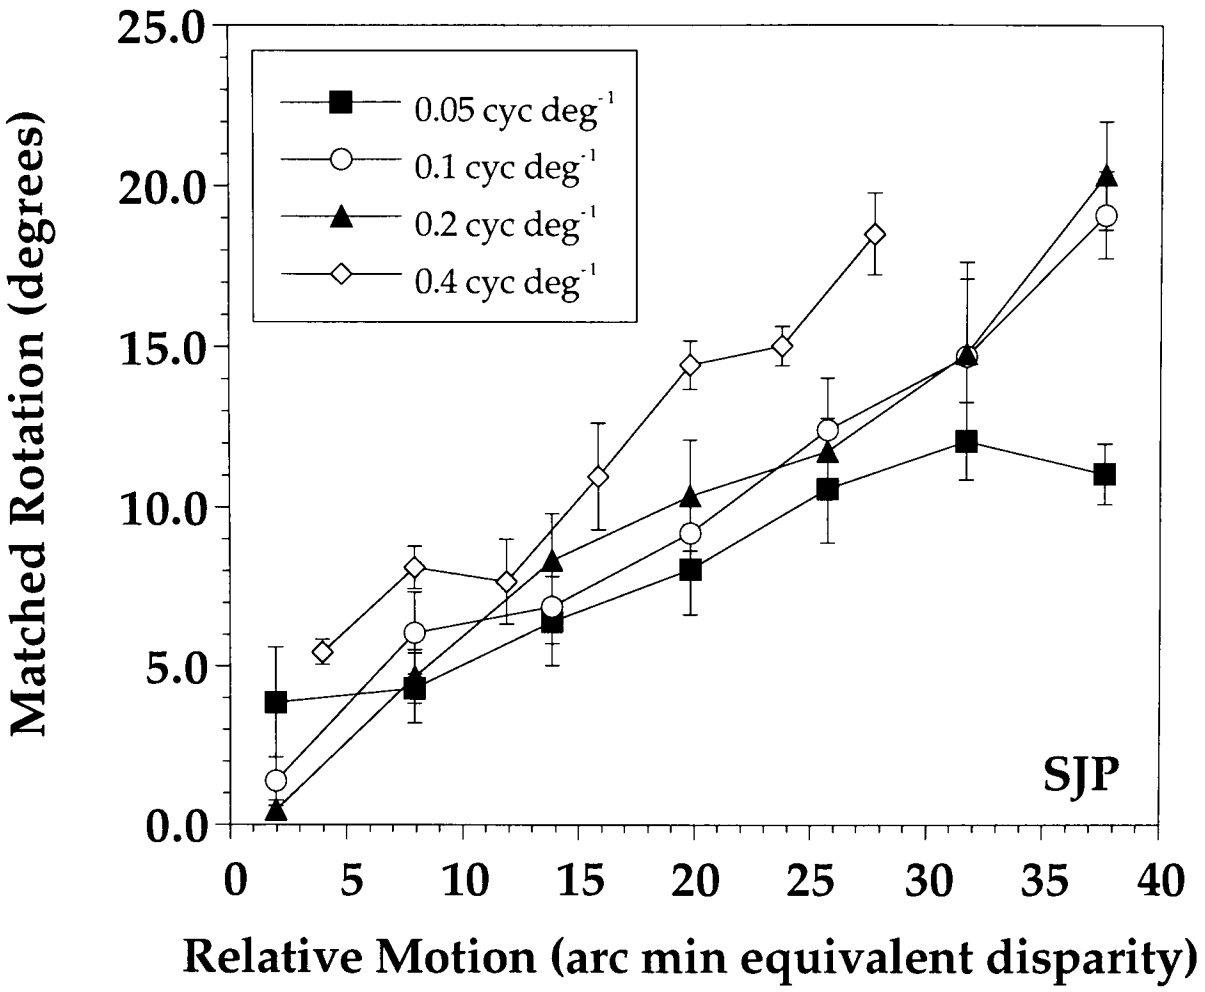
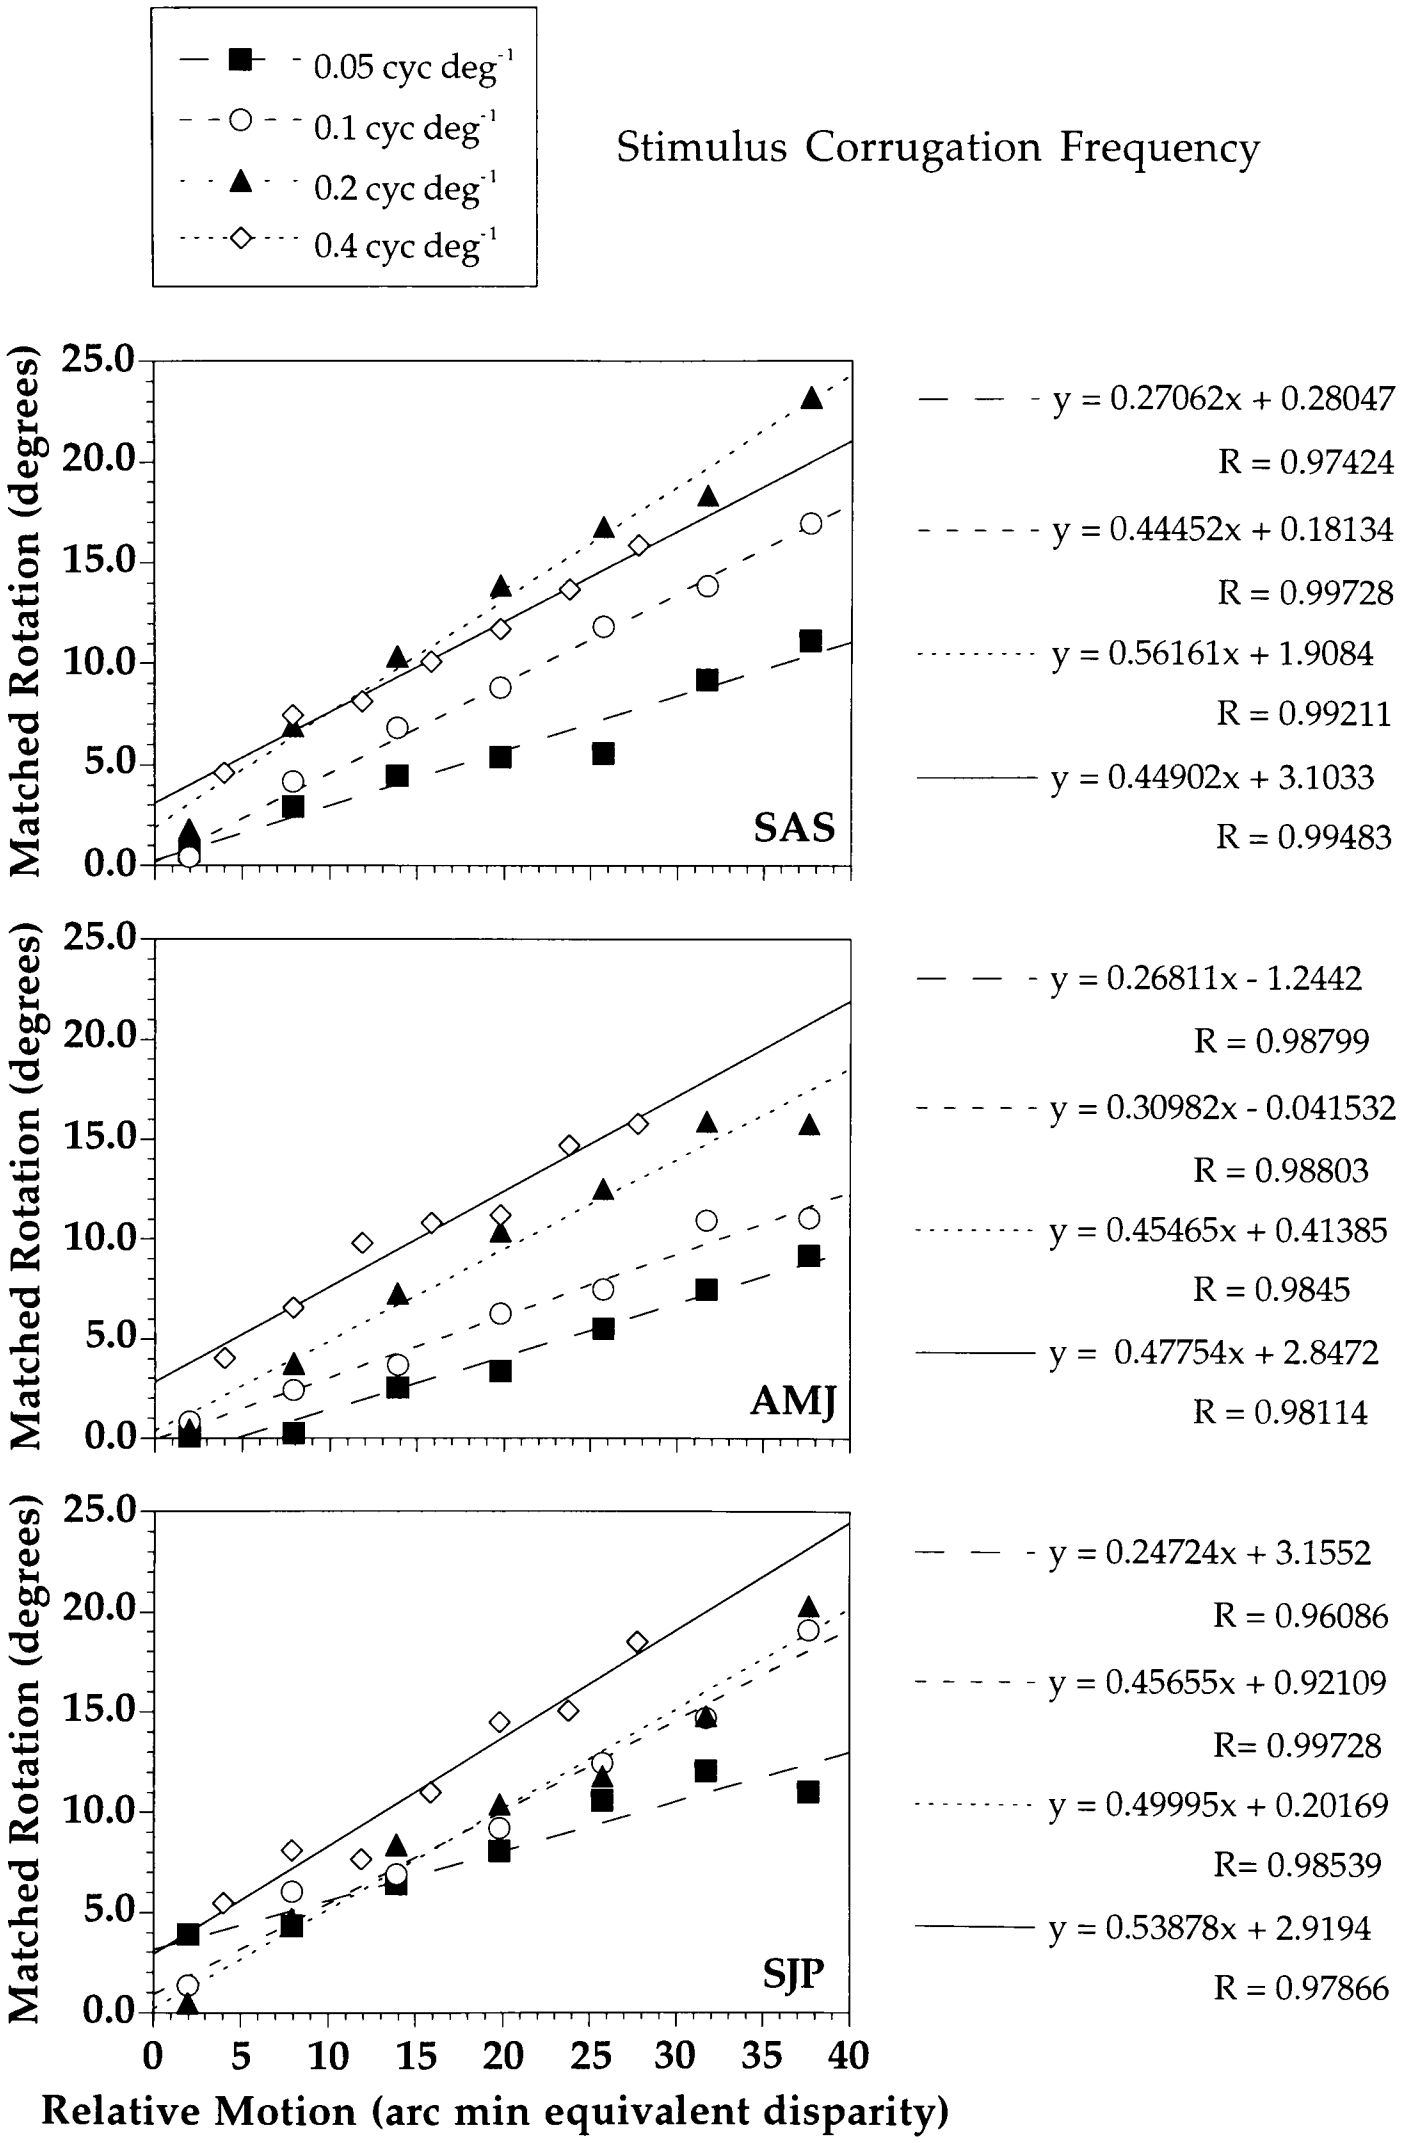


Fig. 4.7: Matched Rotation as a function of Relative Motion for four Corrugation Frequencies for Subject SJP.

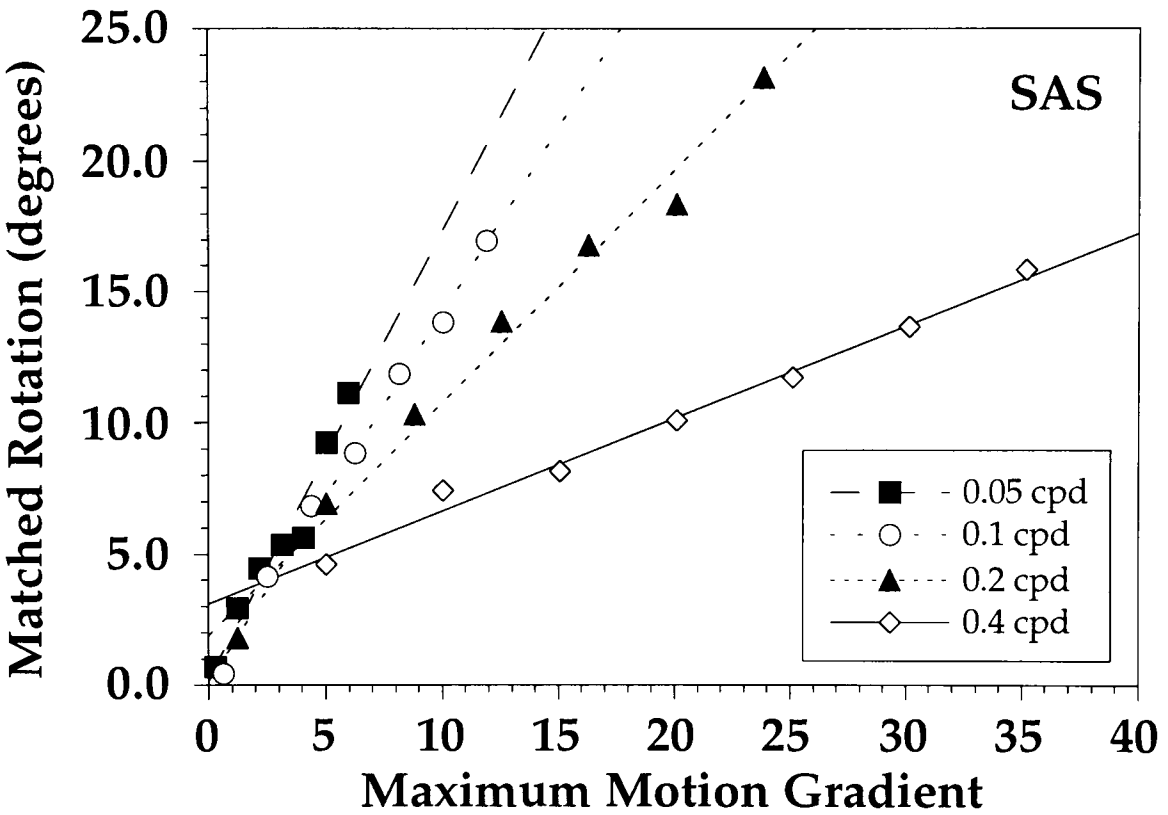
Figs. 4.8 - 4.10 show the lines of best fit for these data, along with their equations.



Figs. 4.8 - 4.10: Matched Rotation as a function of Relative Motion for the four corrugation frequencies, with the lines of best fit and their equations for Subjects SAS, AMJ and SJP.

For a given amount of relative motion, an increase in corrugation frequency generally led to an increased amount of perceived rotation. The slope of the rotation match function generally increased with corrugation frequency. It will be noted that for Subject SAS, however, a greater amount of rotation was perceived in the 0.2 cpd surfaces than in the 0.4 cpd surfaces. This is discussed shortly.

Figs. 4.11 - 4.13 replot the data in terms of the MMG. It can be seen that the slope of the function relating matched rotation to MMG increased with decreasing corrugation frequency.



— — - y = 0.28047 + 1.7063x R= 0.97424

· · · · · y = 0.18134 + 1.4014x R= 0.99728

····· y = 1.9084 + 0.88527x R= 0.99211

——— y = 3.1033 + 0.3539x R= 0.99483

Fig. 4.11: Matched rotation as a function of MMG for Subject SAS.

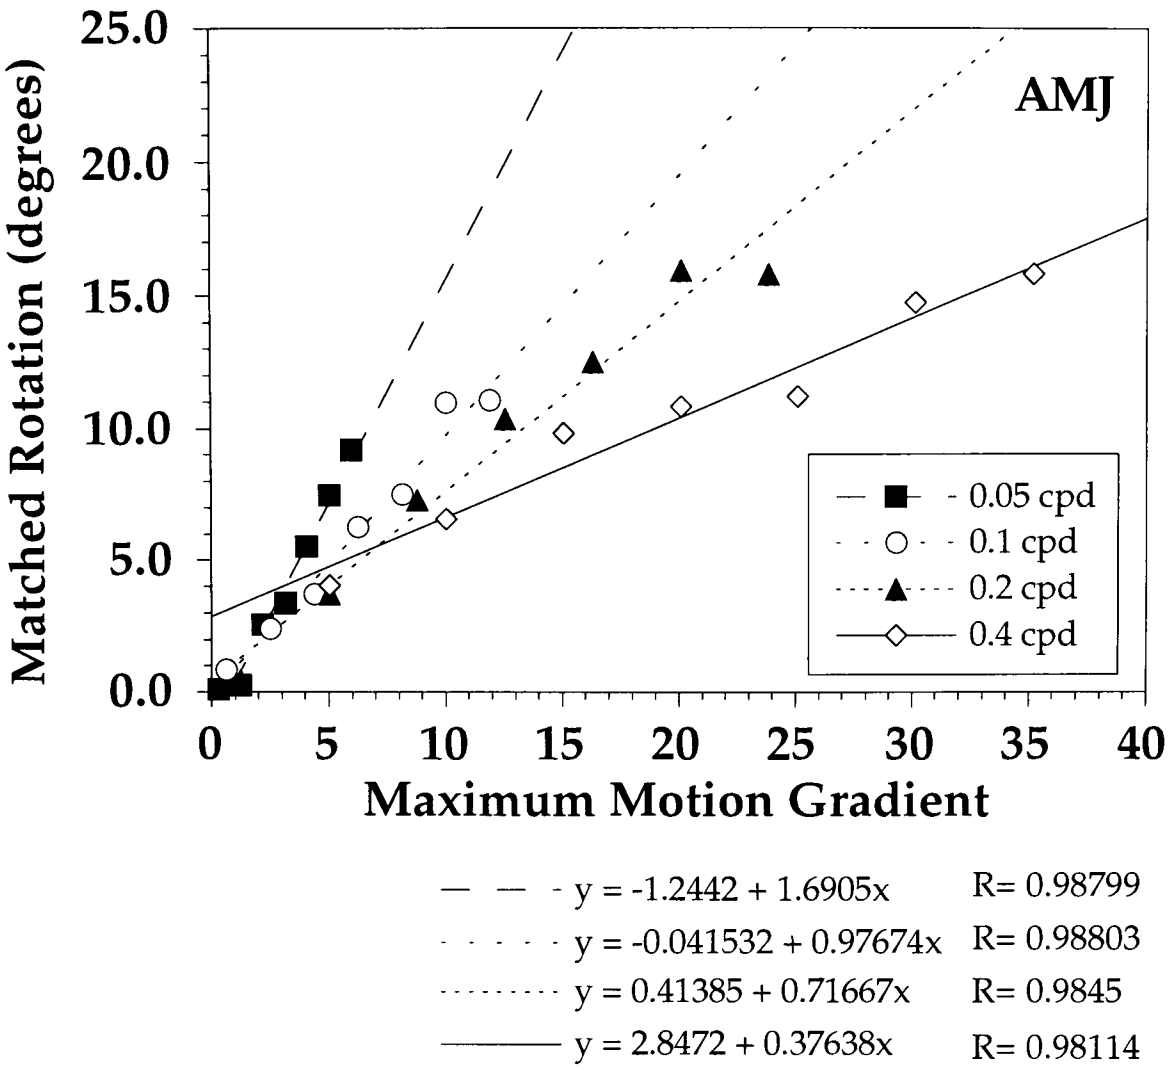


Fig. 4.12: Matched rotation as a function of MMG for Subject AMJ.

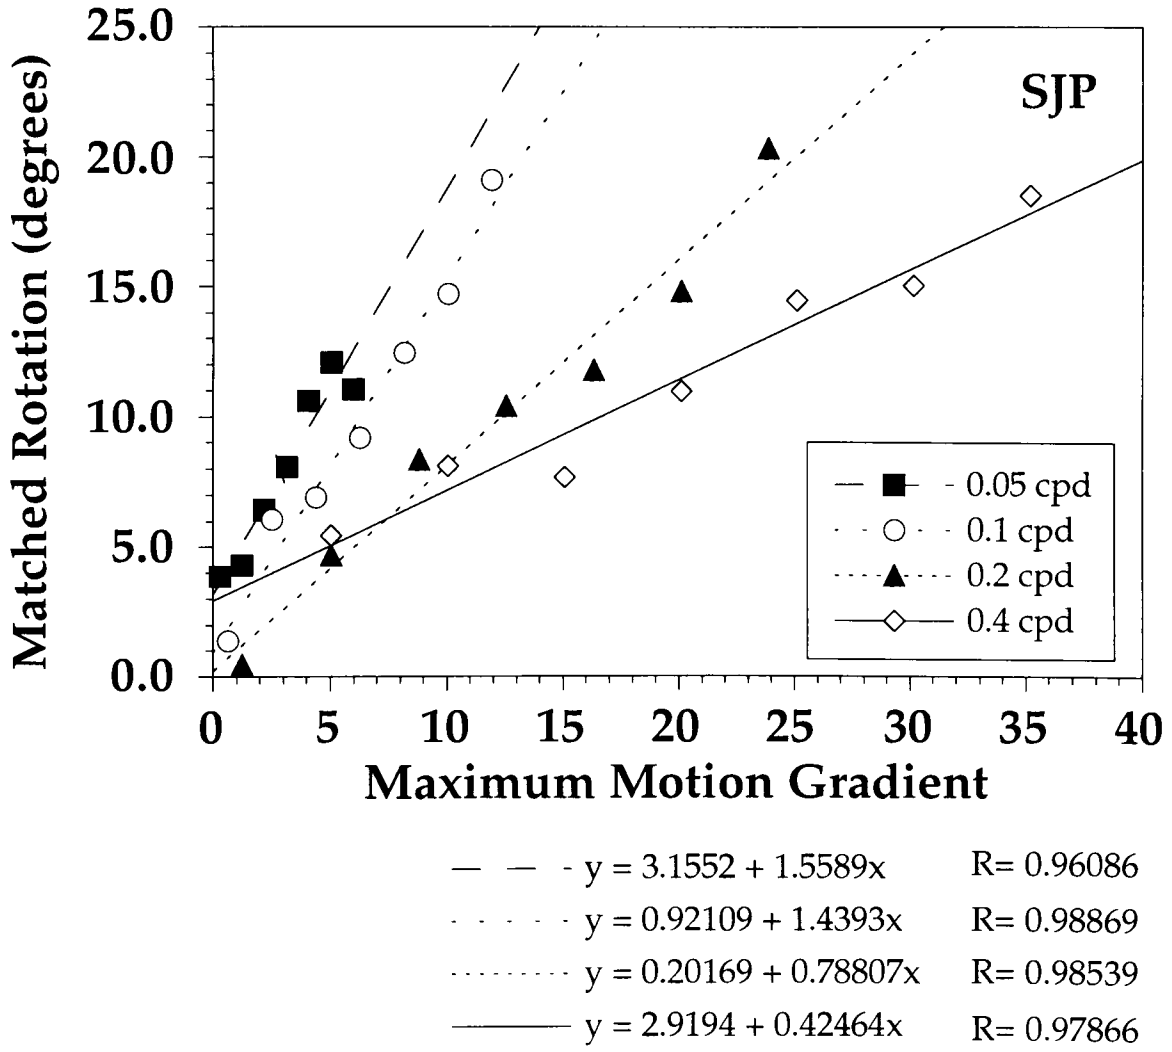


Fig. 4.13: Matched rotation as a function of MMG for Subject SJP.

4.3.5 Discussion

Figs. 4.5 - 4.7 show that for a given amount of relative motion an increase in the stimulus corrugation frequency generally led to an increase in the rotation match. For Subject SAS, the amount of perceived rotation of the 0.4 cpd surface was less than that of the 0.2 cpd surface, however. This point will be addressed shortly. The observed increase in perceived rotation with increasing corrugation frequency complements the finding that Transition Points decrease with increasing corrugation frequency (Experiments 1a and 1b), and with the finding that for a given amount of relative motion, stimuli of high corrugation frequency appear to rotate to a greater extent than stimuli of low corrugation frequency (Experiment 6). It is evident though, that the conclusion drawn in the former experiment, namely that over a four-octave range (three-octave for one subject) the Transition Point was primarily determined by the MMG present in the stimulus, *cannot* be used to predict the dependence of perceived rotation on corrugation frequency for relative motion values greater than the Transition Point.

If perceived rotation were dependent on corrugation frequency, then it would double with a doubling of corrugation frequency for a given value of relative motion. The slopes of matched rotation as a function of relative motion amplitude would also double with a doubling of corrugation frequency. Neither of these predictions is borne out in the data. For a given amount of relative motion, a doubling of corrugation frequency led to an increase in perceived rotation, but not to the extent predicted above. The slopes of the matched rotation functions generally increase with increases in corrugation frequency, but at a lower rate than predicted by a MMG-dependent function.

For Subject SAS, the matched rotation function of the 0.4 cpd surface lay between the functions of the 0.1 and 0.2 cpd stimuli. This may have been due to a portion of the relative motion being perceived as stimulus non-rigidity. The fact that, for a given amount of relative motion, the results of Experiment 6 implied that the 0.4 cpd surface appeared to rotate *more* than the 0.2 cpd surface, while in the current experiment the relative rotation magnitudes were reversed, may be attributable either to the difference between the tasks or to the different relative motion values

used in each experiment: the values of Experiment 6 were generally lower than those of the current experiment, and the lower apparent rotation of the 0.4 cpd surface occurred to a greater extent at higher values of relative motion in the current experiment.

Williams (1993) carried out two experiments involving the measurement of perceived magnitude of rotation in monocularly-viewed motion parallax surfaces. Experiment 1 (p. 81) measured the perceived rotation in parallax surfaces which subtended 8.25° (horizontally) and 6.5° (vertically) and translated horizontally through a distance of 6.5 cm, 6.5° at the 57 cm viewing distance. The corrugation frequency of the surface was 0.14, 0.28 or 0.56 cpd. Subjects were instructed to rate the magnitude of perceived rotation on a six-point scale. Relative motion ranged from -80 to +80 arc min equivalent disparity. It was found that the amount of perceived convex rotation increased with increasing absolute amounts of parallax. Rated rotation increased to a lesser extent for the surface of lowest corrugation frequency (0.14 cpd) than for the two surfaces of higher corrugation frequencies (0.28 and 0.56 cpd). Pooled across relative motion conditions and subjects, rated rotation increased with increasing corrugation frequency. This is in direct agreement with the results of the current experiment.

Figs. 4.11 - 4.13 show how matched rotation varies with the MMG for each corrugation frequency. The MMG has an increasingly strong effect as corrugation frequency decreases. Interestingly, these graphs show that for a given MMG, the amount of perceived rotation increases with decreasing corrugation frequency.

Taken together, the two sets of graphs indicate that perceived rotation is determined both by the MMG and by the absolute amplitude of relative motion.

It will be recalled from Experiment 5 that the Upper Transition Points lay between relative motion values of 3 and 20 arc min equivalent disparity for Subject SAS, and between 7 and 20 arc min equivalent disparity for Subject SJP. A number of the values presented in the current experiment were obviously considerably higher than this, yet subjects were still able to make rotation matches and did not comment on

perceived non-rigidity of the displays. Furthermore, matched rotation generally only started to fall off at relative motion values greater than 30 arc min equivalent disparity, which suggests either that the relative motion previously attributed (in Experiment 5) to non-rigidity was perceived in the current experiment solely as rotation, or that non-rigidity and rotation were simultaneously perceived.

The latter suggestion is likely to be true for the levels of relative motion greatly in excess of the Upper Transition Point; the results of Subject SAS suggest that under certain circumstances this was indeed the case. It is possible, however, that at lower relative motion levels the surface was generally perceived to be rigid, and the relative motion was interpreted as arising from the rotation of the surface rather than from its non-rigidity. This possibility is strengthened by the observation that matched rotation did not fall off significantly at relative motion values greater than the Upper Transition Points determined in Experiment 5. The fact that there was no change in the slope of the function as the amount of relative motion exceeded this value suggests that the perceptual change which occurred in Experiment 5 (i.e. the change from a rigid, rotating surface to a non-rigid, rotating surface) did not occur in the current experiment.

If non-rigidity was not perceived in the current experiment then its explanation might lie along the following lines. Under certain conditions observers consciously experience a percept of a given nature only when instructed to make judgements concerning it. This is unsurprising given the ambiguous or metameric characteristics of the stimuli used here. Although some stimulus attributes obviously exist whether or not attention is directed towards them, for example perceived depth or rotation, it would seem that other attributes are less readily perceived unless the observers specifically direct their attention towards them in order to meet the requirements of the task. At relatively low parallax amplitudes rigidity appears to be such an attribute. In Experiment 5, subjects' attention was directed towards the apparent rigidity of the parallax surfaces in order that the Upper Transition Points could be determined. In Experiments 6 and 7a, subjects' attention was directed towards apparent rotation. This is not to suggest that non-rigidity was never perceived in these latter experiments, rather that there was perhaps

a decreased likelihood of it being perceived, a notion which accords with observers' general preference of perceiving ambiguous motion as arising from a rigid structure wherever possible.

Motion parallax studies which have investigated the perceived motion of sinusoidally depth modulated surfaces have not reported percepts of non-rigidity, even when the relative motion amplitudes have been considerably higher than the Upper Transition Points determined in Experiment 5. Ono, Rivest and Ono (1986) and Ono and Steinbach (1990) reported that amplitudes of up to 113 arc min equivalent disparity were perceived as rigid surfaces which appeared to some observers to rotate about a vertical axis.

4.4 Experiment 7b: Matched Rotation as a function of Corrugation Frequency of Peak-Clipped surfaces

4.4.1 Introduction

The previous experiment demonstrated that the perceived rotation of surfaces with relative motion amplitudes greater than the Transition Point is affected both by the MMG and the absolute relative motion amplitude. The present experiment used clipped surfaces similar to those of Experiment 1b in order to further investigate the relationship between the MMG and perceived rotation. Experiment 1b found that the Transition Points of surfaces with identical MMGs were generally similar. The current experiment determines whether these surfaces appear to rotate to the same extents at relative motion amplitudes greater than the Transition Point.

4.4.2 Stimuli

Clipped stimuli were created in the same way as described in Section 3.8.2. Stimuli were clipped at 50, 25 or 12.5% of their (unclipped) height. Their MMGs were thus respectively twice, four or eight times that of the unclipped surface. Corrugation frequencies were 0.05, 0.1 and 0.2 cpd. Five values of relative motion were used.

4.4.3 Procedure

This was identical to that of the previous experiment, except that there were five relative motion amplitudes. Trials were blocked in groups of ten, and subjects completed five sessions of one block per corrugation frequency per session.

4.4.4 Results

Figs. 4.14 - 4.21 plot rotation matches for stimuli of the same corrugation frequency but different MMG. These graphs show firstly that the magnitude of perceived rotation increased with increasing relative motion amplitude. Secondly, with increased levels of clipping, perceived rotation generally increased for the 0.05 cpd surface. This trend was not found for either the 0.1 or the 0.2 cpd surface.

Figs. 4.14 and 4.16 show that for the corrugation frequency 0.05 cpd, the matched rotation function of the 12.5%-clipped surface lies above that of the unclipped surface. The rotation match data of the unclipped and the 12.5%-clipped 0.05 cpd surfaces are shown separately in Figs. 4.15 and 4.17 for Subjects SAS and SJP respectively, in order to make the relative positions of these two functions clearer. The functions of the 25%- and 50%-clipped surfaces lie between these. Figs. 4.18 - 4.21 show that variation in the MMG of 0.1 and 0.2 cpd stimuli has little effect on the matched rotation function.

Two-way repeated-measures ANOVAs were carried out on the data of each subject, with clippedness and relative motion amplitude as the factors, and trial repetition as the repeated factor. Separate ANOVAs were carried out for each corrugation frequency since the number of levels of clippedness varied with the former variable.

For the 0.05 cpd surface, there was a main effect of clippedness for each subject ($F[3, 24^{\dagger}] = 9.62$, $p<0.0005$; $F[3, 27] = 11.78$, $p<0.0005$, for Subjects SAS and SJP respectively). (†NOTE: the degrees of freedom of the error term should theoretically be 27 in both cases ([clippedness levels - 1] * [number of repeated trials - 1]). However, an error during the experiment resulted in nine (as opposed to ten) repetitions for the 0.05 cpd 50%-clipped and

0.05 cpd 25%-clipped surfaces for Subject SAS, resulting in the observed 24 degrees of freedom. This is the explanation for the differences in the degrees of freedom of the error terms for the two subjects in the remainder of the 0.05 cpd data analyses). There was a main effect of relative motion amplitude ($F[4, 32] = 82.10$, $p<0.0005$; $F[4, 36] = 160.54$, $p<0.0005$, for Subjects SAS and SJP respectively). There was a significant interaction between clippedness and relative motion amplitude ($F[12, 96] = 1.95$, $p<0.05$; $F[12, 108] = 3.51$, $p<0.0005$, for Subjects SAS and SJP respectively).

For the 0.1 cpd surface, there was no main effect of clippedness for either subject ($F[2, 18] = 0.67$, n.s.; $F[2, 18] = 0.47$, n.s., for Subjects SAS and SJP respectively). There was a main effect of relative motion amplitude for both subjects ($F[4, 36] = 171.03$, $p<0.0005$; $F[4, 36] = 235.95$, $p<0.0005$, for Subjects SAS and SJP respectively). The interaction between clippedness and relative motion amplitude was non-significant for both subjects ($F[8, 72] = 0.45$, n.s.; $F[8, 72] = 1.91$, n.s., for Subjects SAS and SJP respectively).

For the 0.2 cpd surface, there was no main effect of clippedness for either subject ($F[1, 9] = 0.59$, n.s.; $F[1, 9] = 91.83$, n.s., for Subjects SAS and SJP respectively). There was a main effect of relative motion amplitude for both subjects ($F[4, 36] = 110.68$, $p<0.0005$; $F[4, 36] = 221.83$, $p<0.0005$, for Subjects SAS and SJP respectively). The interaction between clippedness and relative motion amplitude was non-significant for both subjects ($F[4, 36] = 0.48$, n.s.; $F[4, 36] = 0.29$, n.s., for Subjects SAS and SJP respectively).

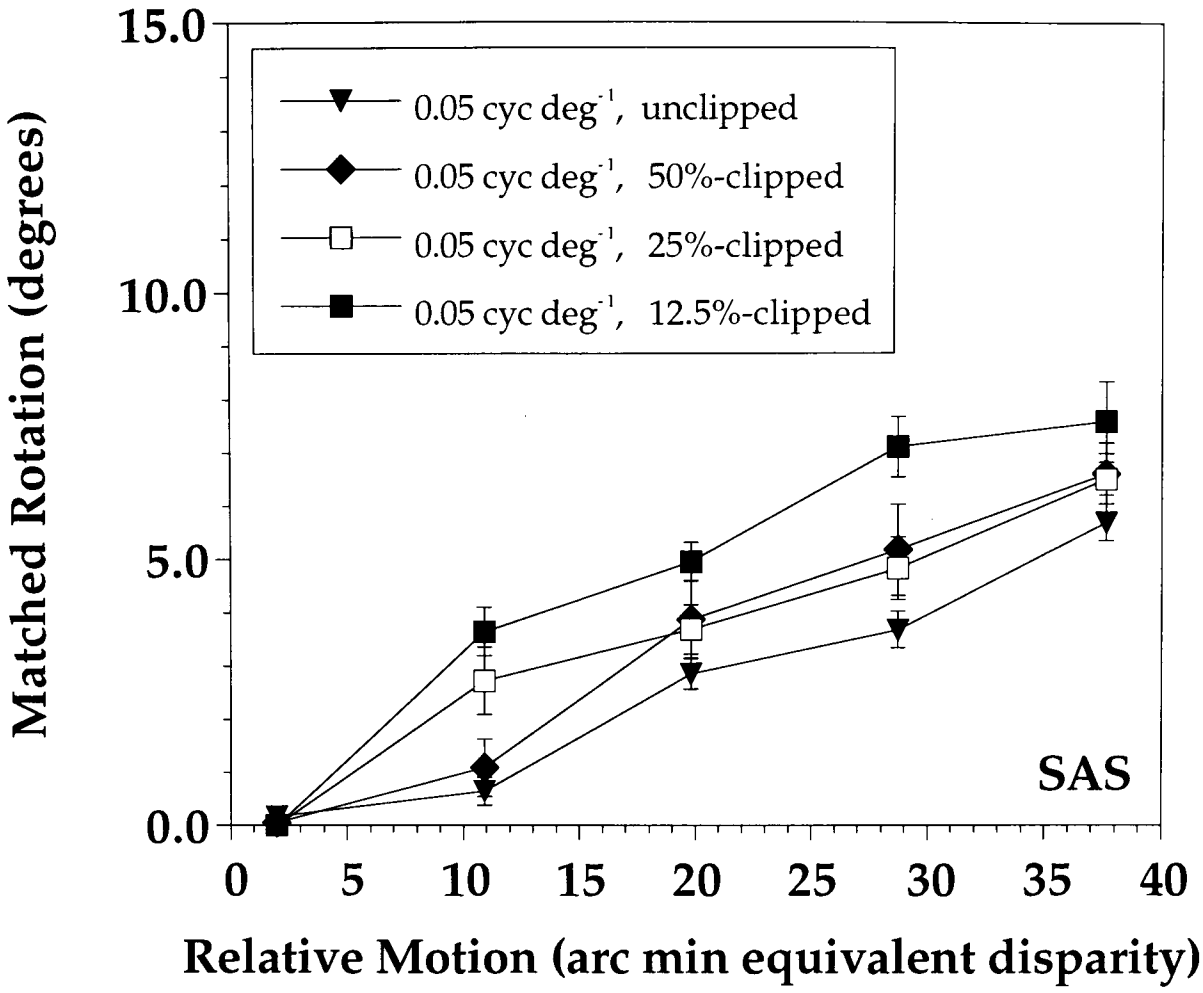


Fig. 4.14: Matched Rotation as a function of Relative Motion for a 0.05 cpd surface clipped to varying degrees, for Subject SAS.

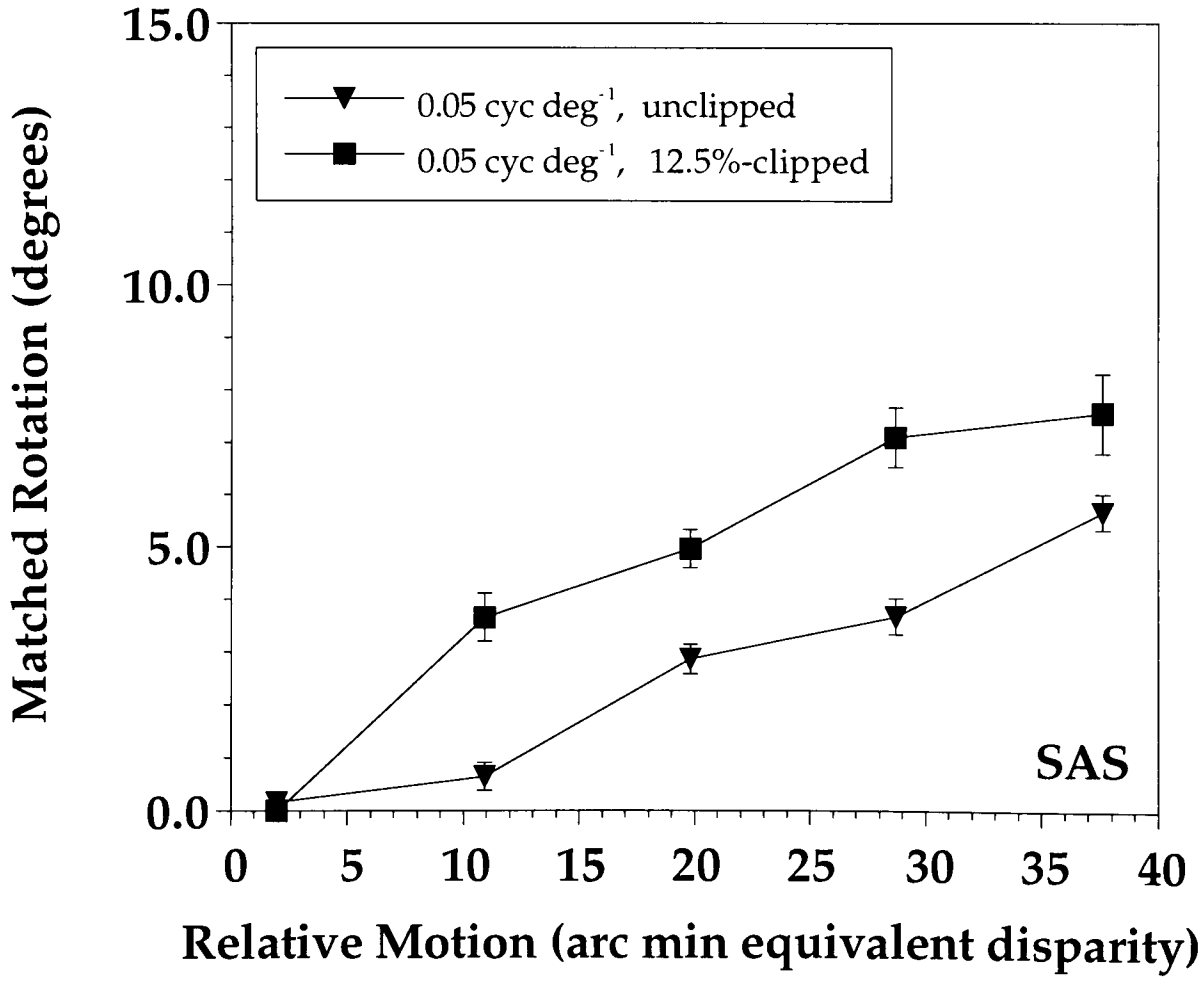


Fig. 4.15: Matched Rotation as a function of Relative Motion for a 0.05 cpd surface (unclipped and 12.5%-clipped conditions only), for Subject SAS.

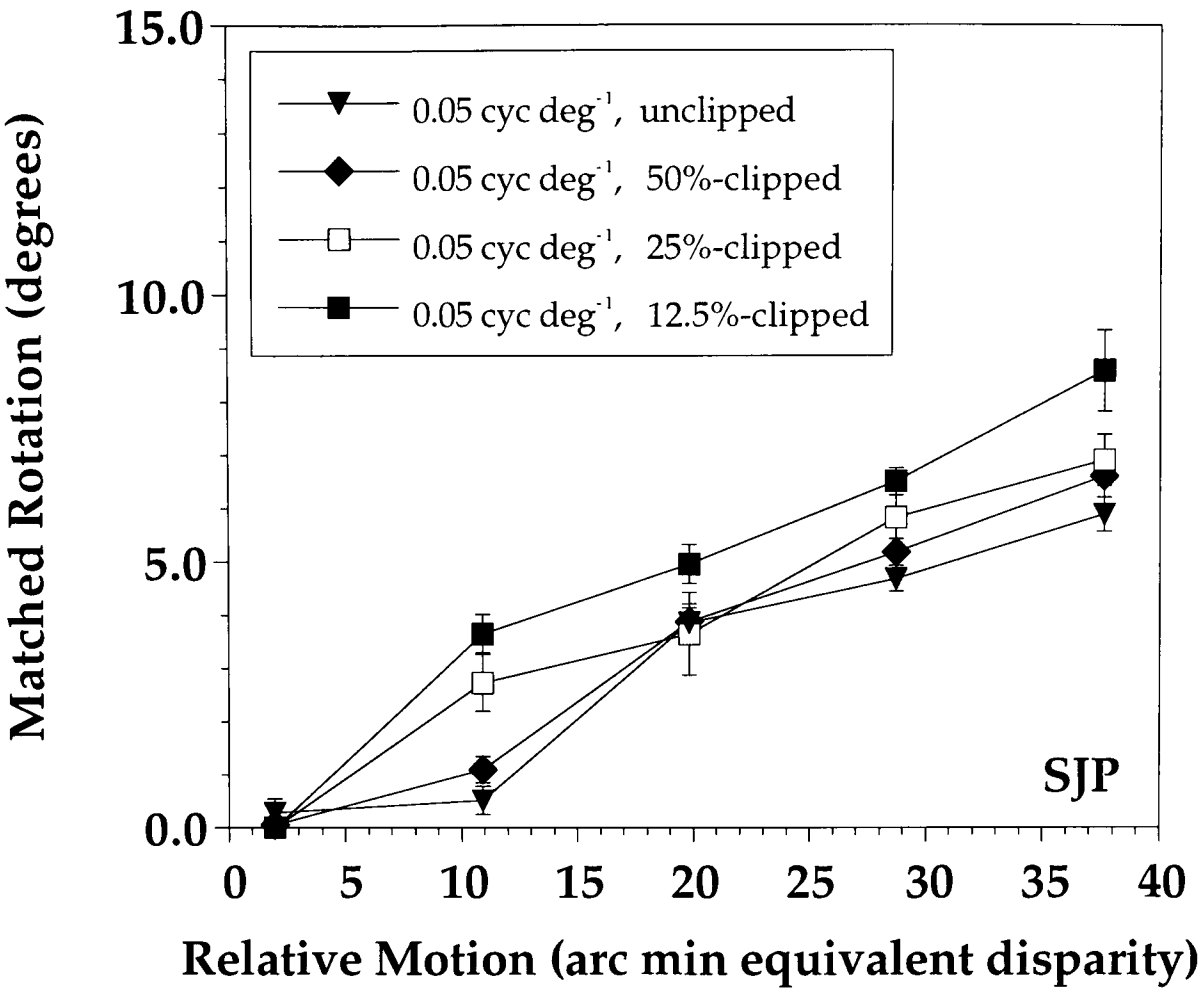


Fig. 4.16: Matched Rotation as a function of Relative Motion for a 0.05 cpd surface clipped to varying degrees, for Subject SJP.

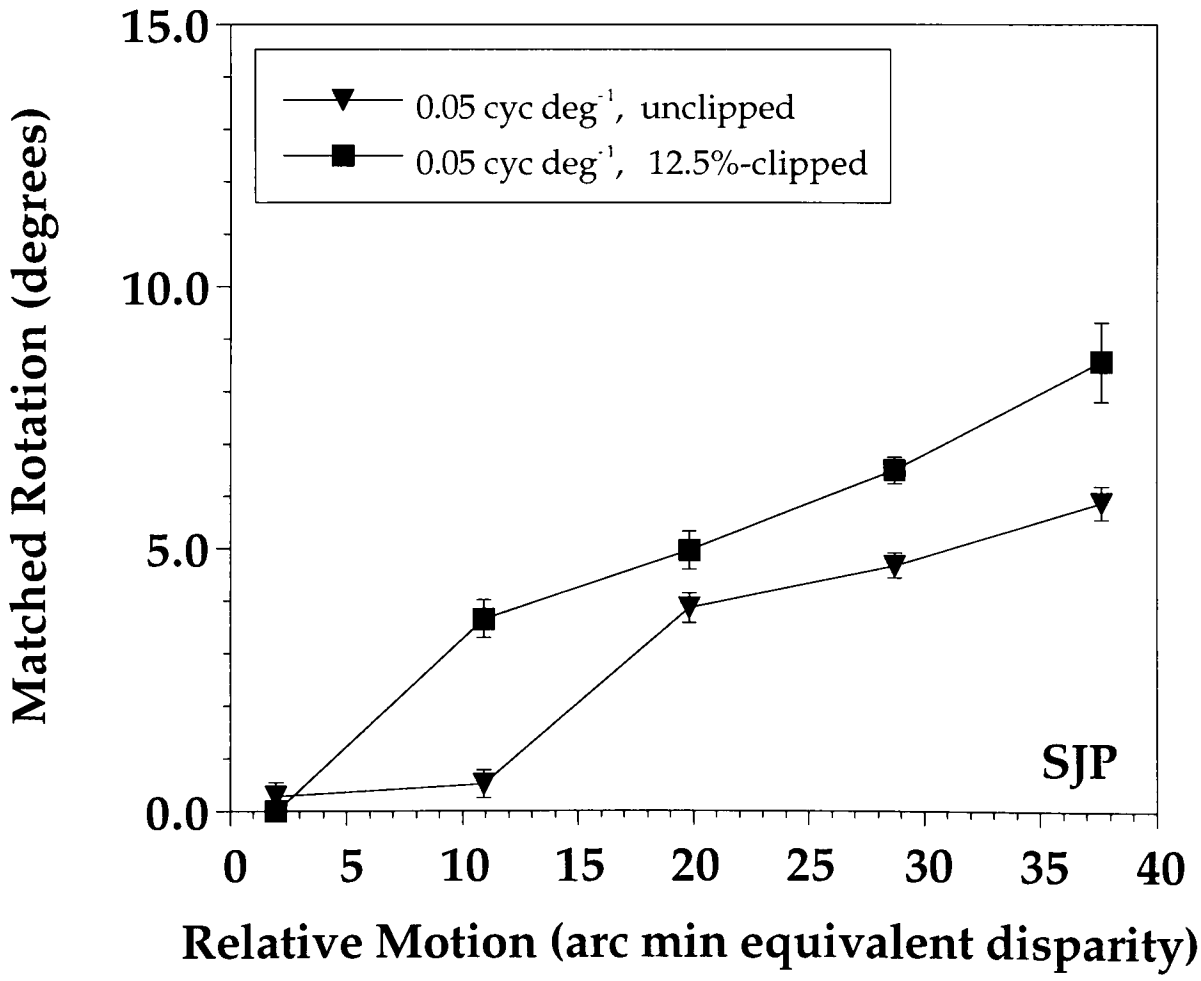


Fig. 4.17: Matched Rotation as a function of Relative Motion for a 0.05 cpd surface (unclipped and 12.5%-clipped conditions only), for Subject SJP.

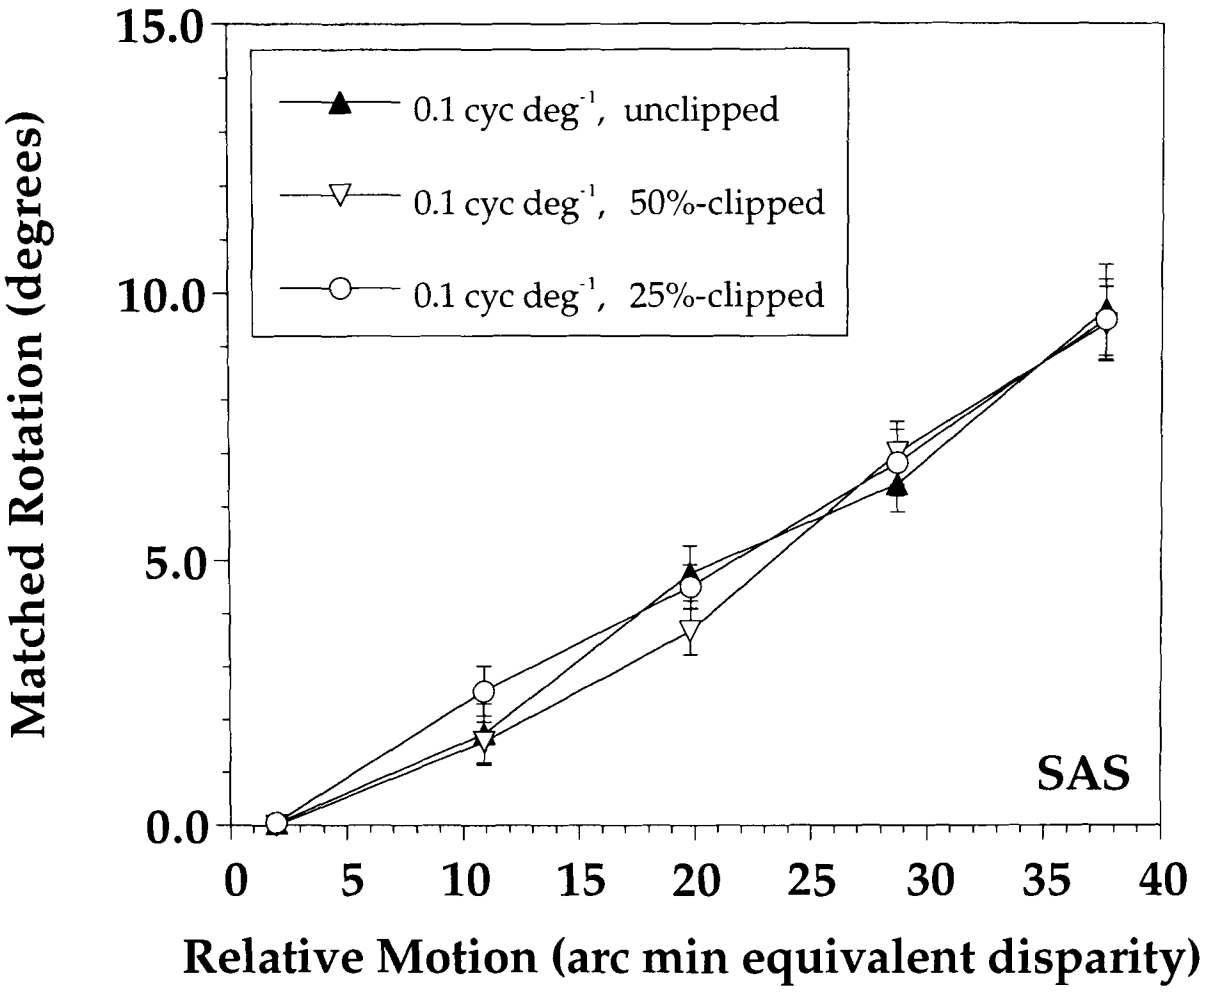


Fig. 4.18: Matched Rotation as a function of Relative Motion for a 0.1 cpd surface clipped to varying degrees, for Subject SAS.

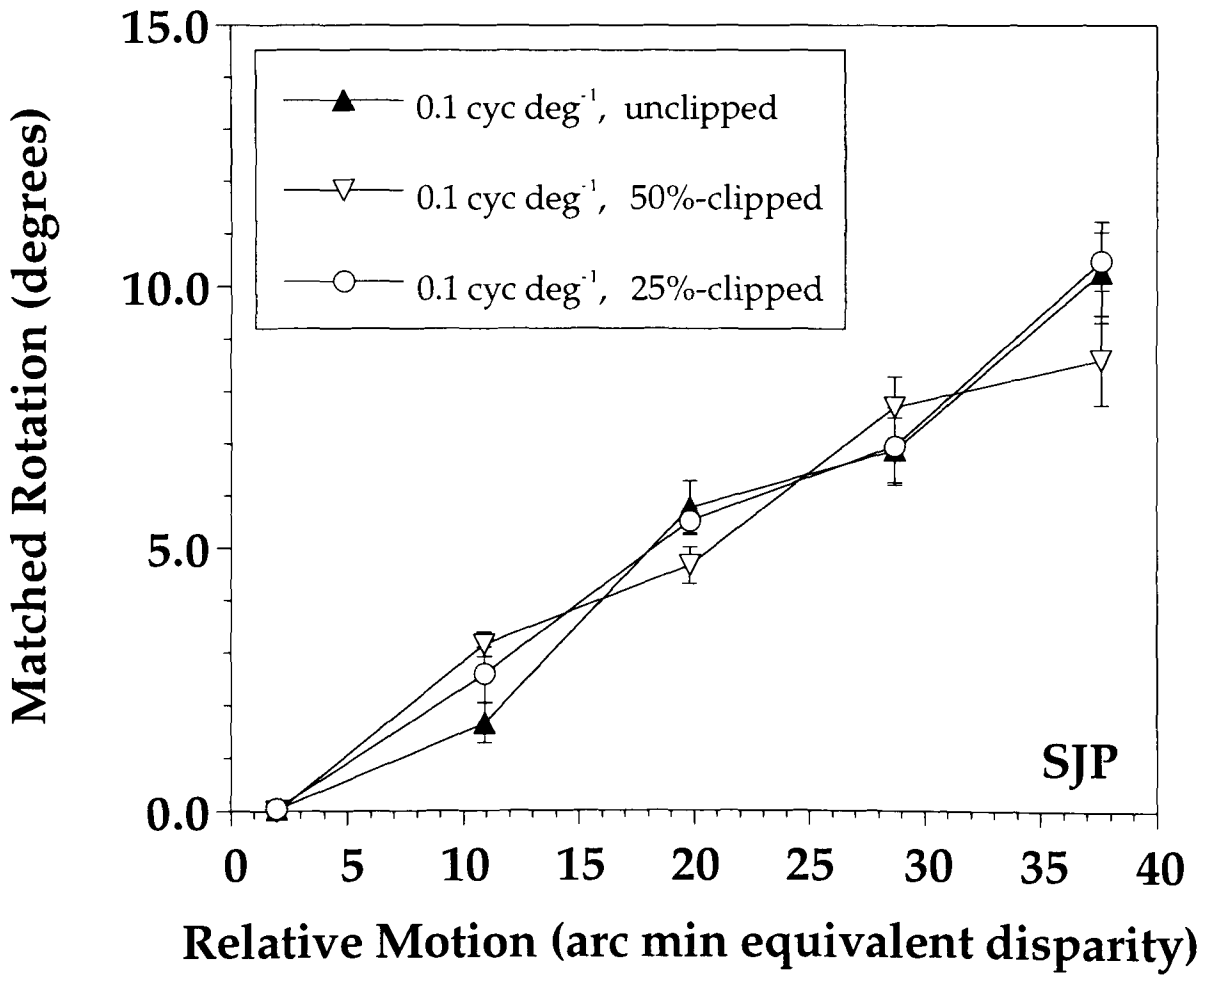


Fig. 4.19: Matched Rotation as a function of Relative Motion for a 0.1 cpd surface clipped to varying degrees, for Subject SJP.

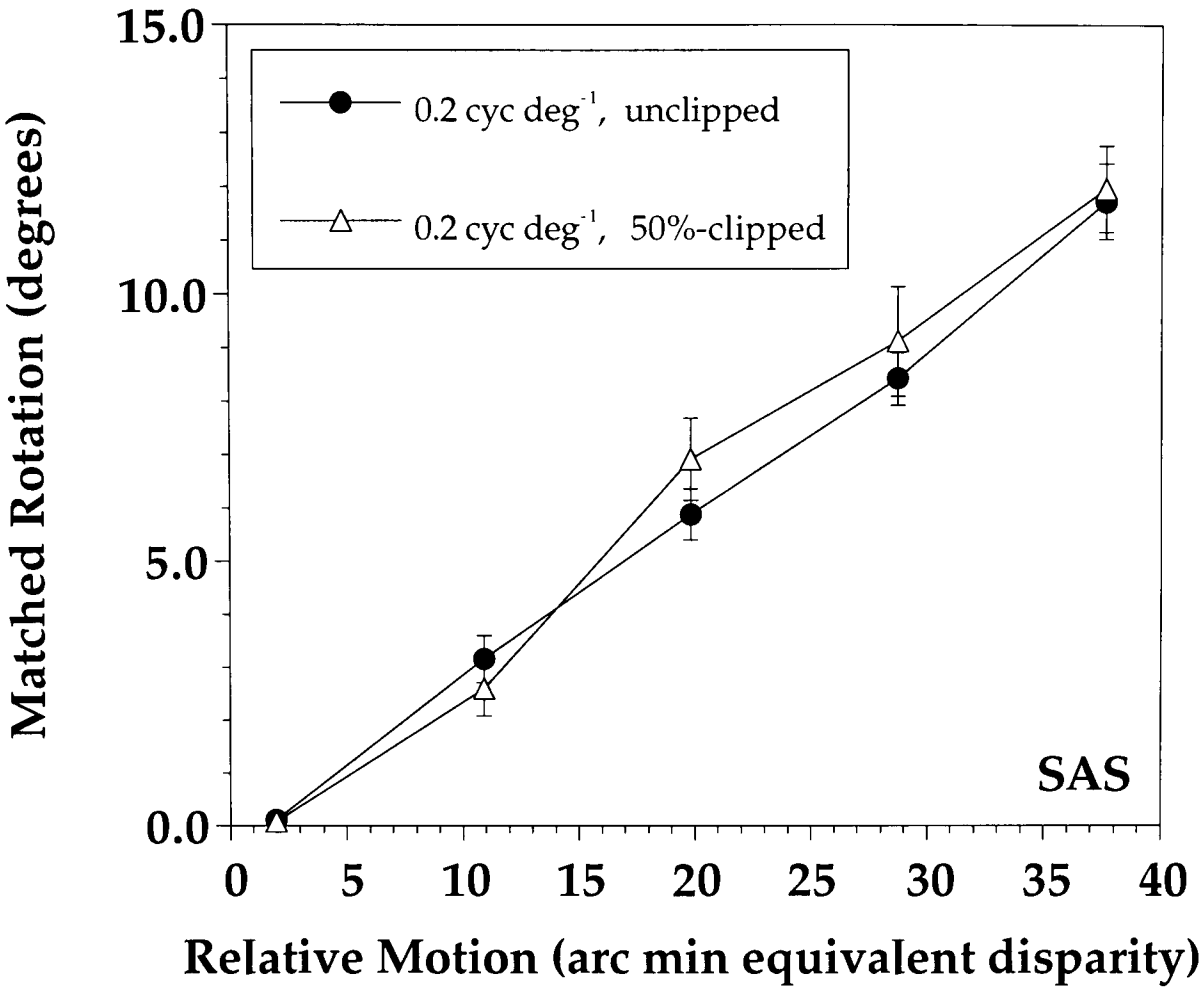


Fig. 4.20: Matched Rotation as a function of Relative Motion for a 0.2 cpd surface clipped to varying degrees, for Subject SAS.

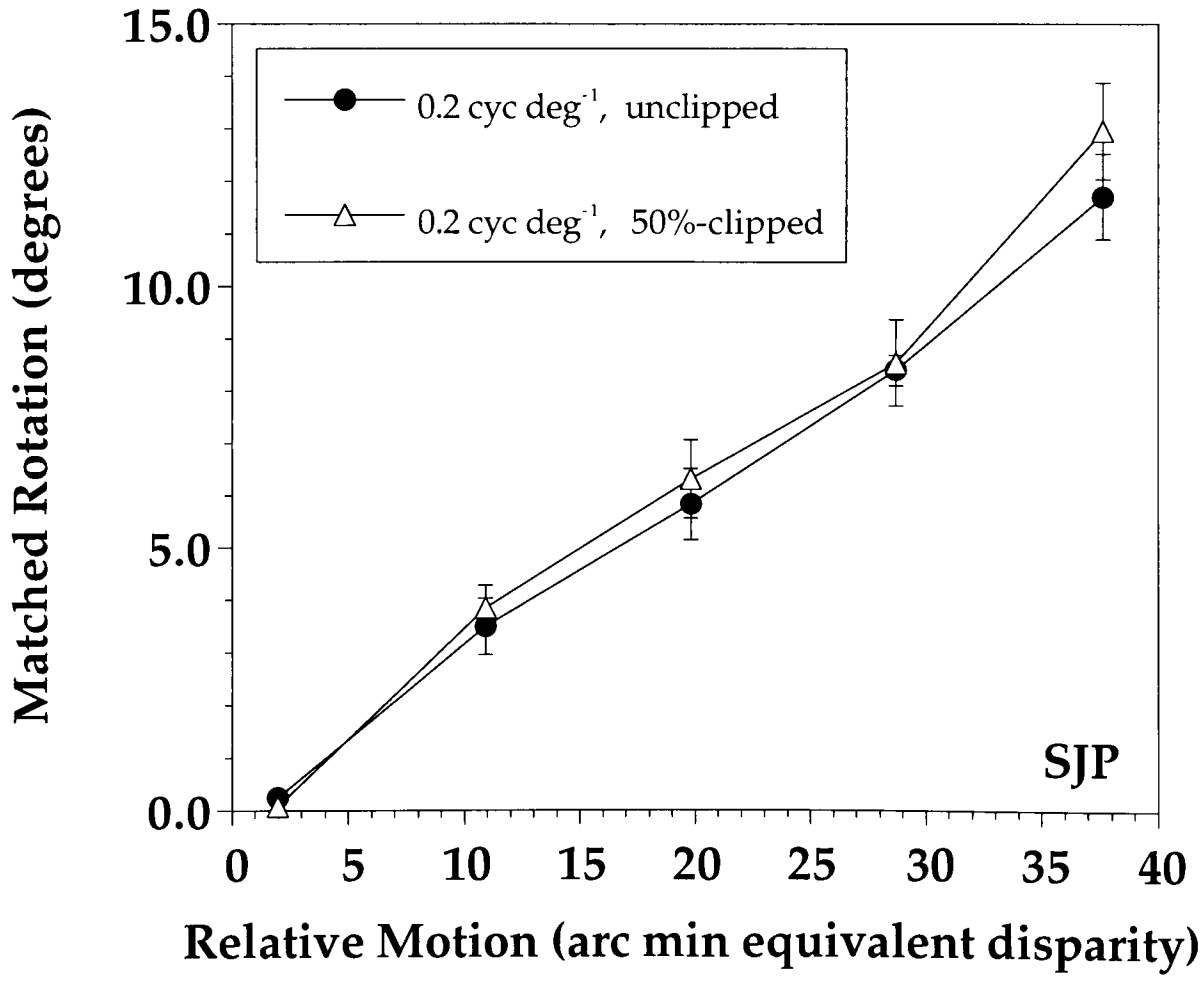


Fig. 4.21: Matched Rotation as a function of Relative Motion for a 0.2 cpd surface clipped to varying degrees, for Subject SJP.

The data are replotted in Figs. 4.22 - 4.29 to show rotation matches for stimuli of different corrugation frequencies but equal MMGs. Once more, for the sake of clarity, Figs. 4.23 and 4.25 show the rotation match data of the 0.05 cpd 12.5%-clipped and the 0.4 cpd unclipped surfaces alone, for Subjects SAS and SJP respectively.

Surfaces of equal MMGs did not produce equal rotation matches. For a given amount of relative motion, matches increased with increasing corrugation frequency.

Two-way repeated-measures ANOVAs were carried out on the data of each subject, with corrugation frequency and relative motion amplitude as the factors, and trial repetition as the repeated factor. Separate ANOVAs were carried out for each set of surfaces with identical MMGs, since the number of levels of corrugation frequency varied with the former variable.

For surfaces with MMGs equal to that of a 0.4 cpd unclipped surface (i.e. those of Fig. 4.22 and 4.24) there was a main effect of corrugation frequency for both subjects ($F[3, 27] = 9.57$, $p<0.0005$; $F[3, 27] = 15.77$, $p<0.0005$, for Subjects SAS and SJP respectively). There was a main effect of relative motion amplitude for both subjects ($F[4, 36] = 160.68$, $p<0.0005$; $F[4, 36] = 233.31$, $p<0.0005$, for Subjects SAS and SJP respectively). There was a significant interaction between corrugation frequency and relative motion amplitude for both subjects ($F[12, 108] = 3.23$, $p<0.05$; $F[12, 108] = 3.86$, $p<0.0005$, for Subjects SAS and SJP respectively).

For surfaces with MMGs equal to that of a 0.2 cpd unclipped surface (i.e. those of Fig. 4.26 and 4.27) there was a main effect of corrugation frequency for both subjects ($F[2, 16^{\dagger}] = 26.43$, $p<0.0005$; $F[2, 18] = 14.51$, $p<0.0005$, for Subjects SAS and SJP respectively). (†NOTE the explanation of the differing degrees of freedom of the error term for the two subjects is given in Section 4.4.4 above). There was a main effect of relative motion amplitude for both subjects ($F[4, 32] = 187.82$, $p<0.0005$; $F[4, 36] = 174.03$, $p<0.0005$, for Subjects SAS and SJP respectively). There was a significant interaction between corrugation frequency and relative motion amplitude for both subjects ($F[8, 64] = 4.52$, $p<0.0005$; $F[8, 72] = 2.96$, $p<0.01$, for Subjects SAS and SJP respectively).

For surfaces with MMGs equal to that of a 0.1 cpd unclipped surface (i.e. those of Fig. 4.28 and 4.29) there was a main effect of corrugation frequency for both subjects ($F[1, 8] = 11.30$, $p<0.05$; $F[1, 9] = 21.29$, $p<0.01$, for Subjects SAS and SJP respectively). There was a main effect of relative motion amplitude for both subjects ($F[4, 32] = 59.42$, $p<0.0005$; $F[4, 36] = 139.05$, $p<0.0005$, for Subjects SAS and SJP respectively). There was a significant interaction between corrugation frequency and relative motion amplitude for both subjects ($F[4, 32] = 3.44$, $p<0.05$; $F[4, 36] = 4.68$, $p<0.01$, for Subjects SAS and SJP respectively).

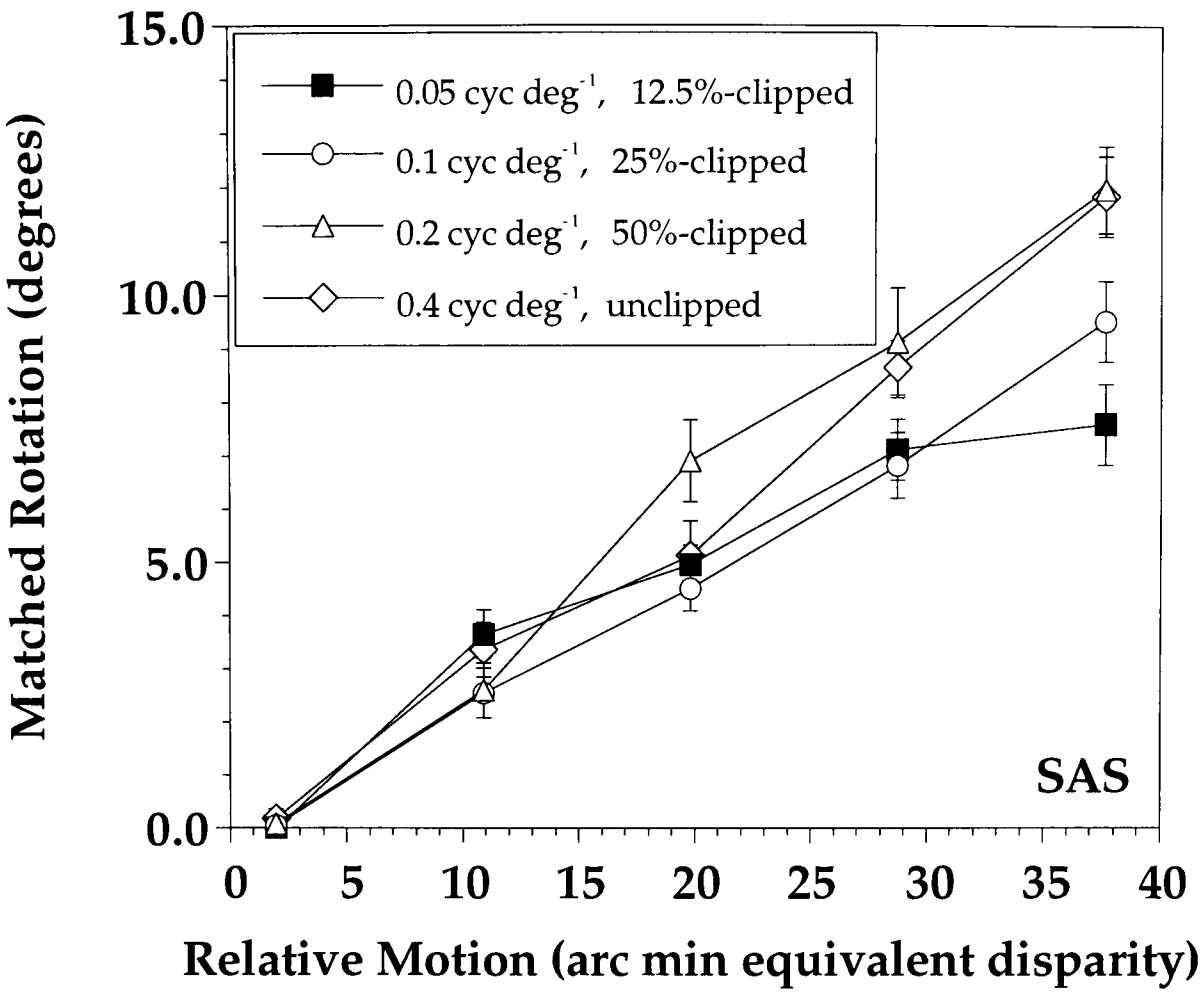


Fig. 4.22: Matched Rotation as a function of Relative Motion for surfaces containing a MMG identical to that of an unclipped 0.4 cpd surface, for Subject SAS.

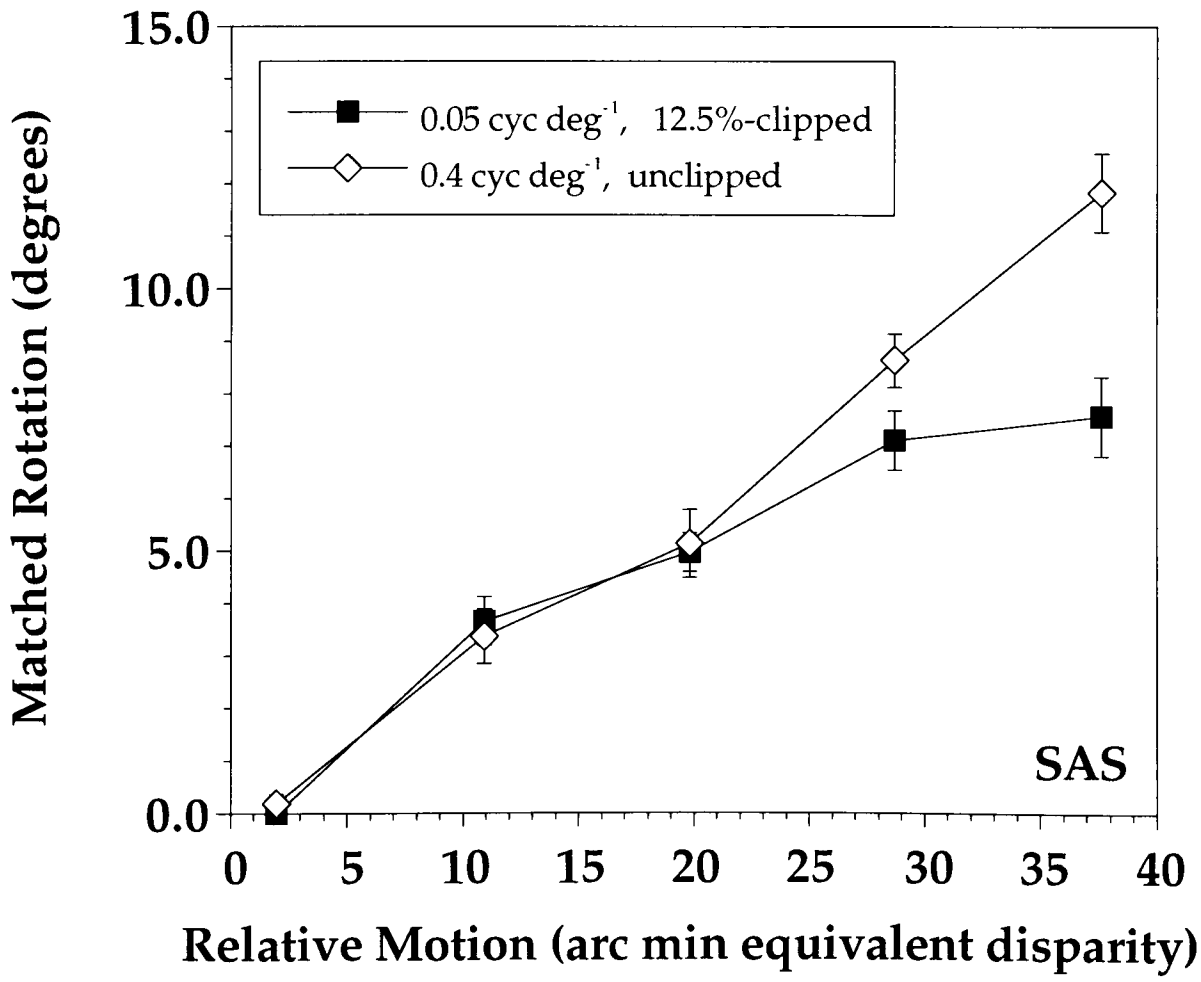


Fig. 4.23: Matched Rotation as a function of Relative Motion for unclipped 0.4 cpd and 12.5%-clipped 0.05 cpd surfaces, for Subject SAS.

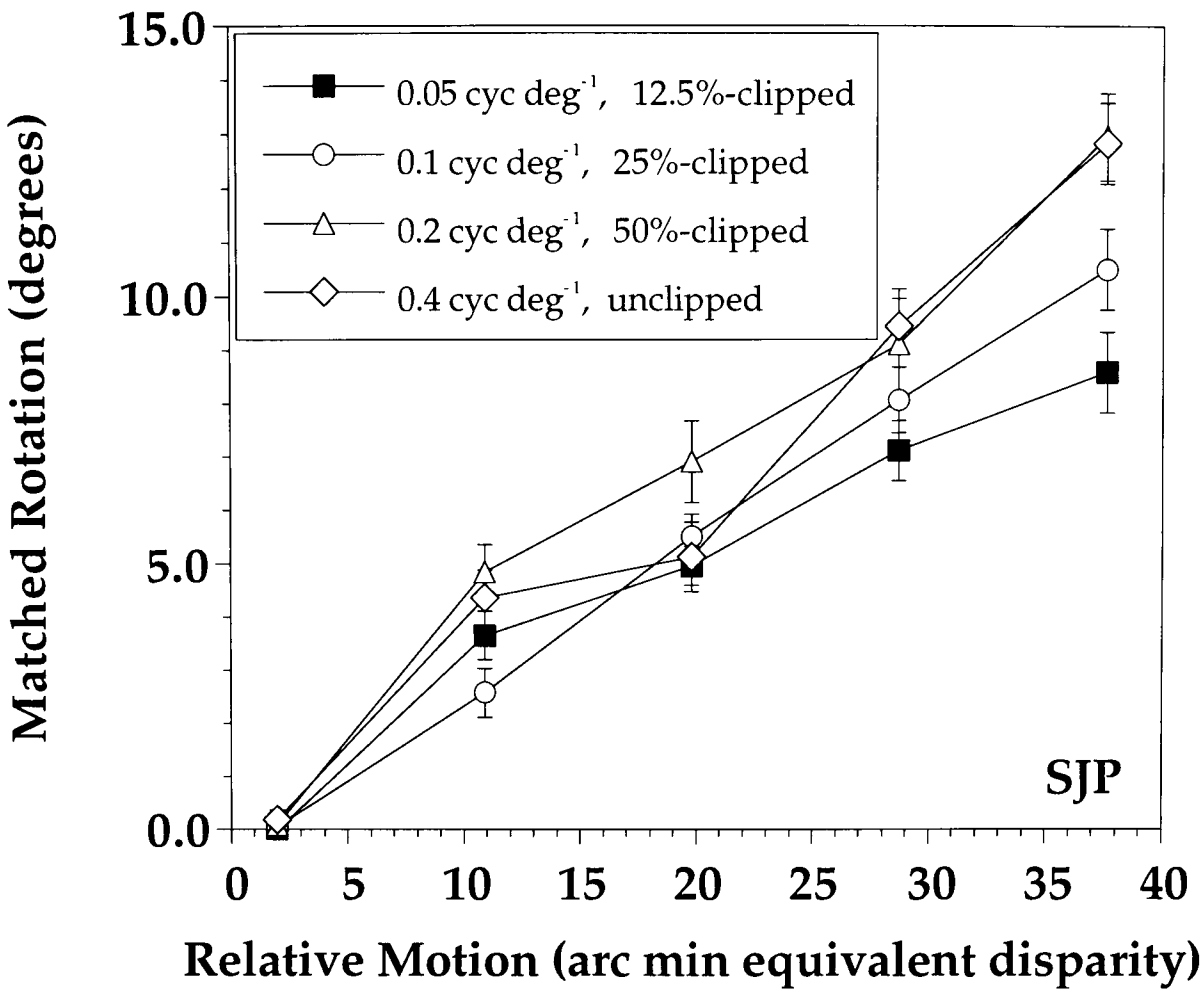


Fig. 4.24: Matched Rotation as a function of Relative Motion for surfaces containing a MMG identical to that of an unclipped 0.4 cpd surface, for Subject SJP.

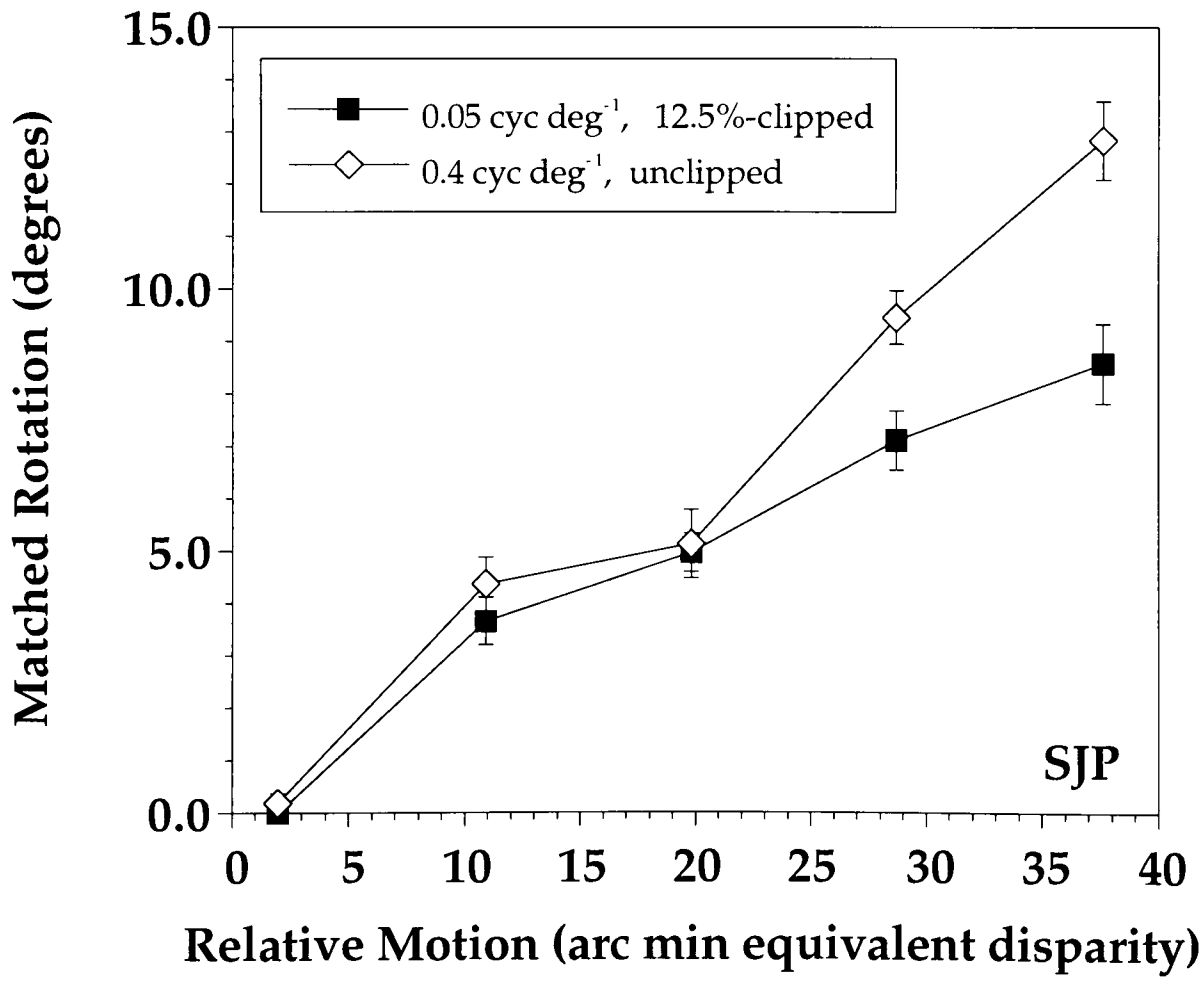


Fig. 4.25: Matched Rotation as a function of Relative Motion for unclipped 0.4 cpd and 12.5%-clipped 0.05 cpd surfaces, for Subject SJP.

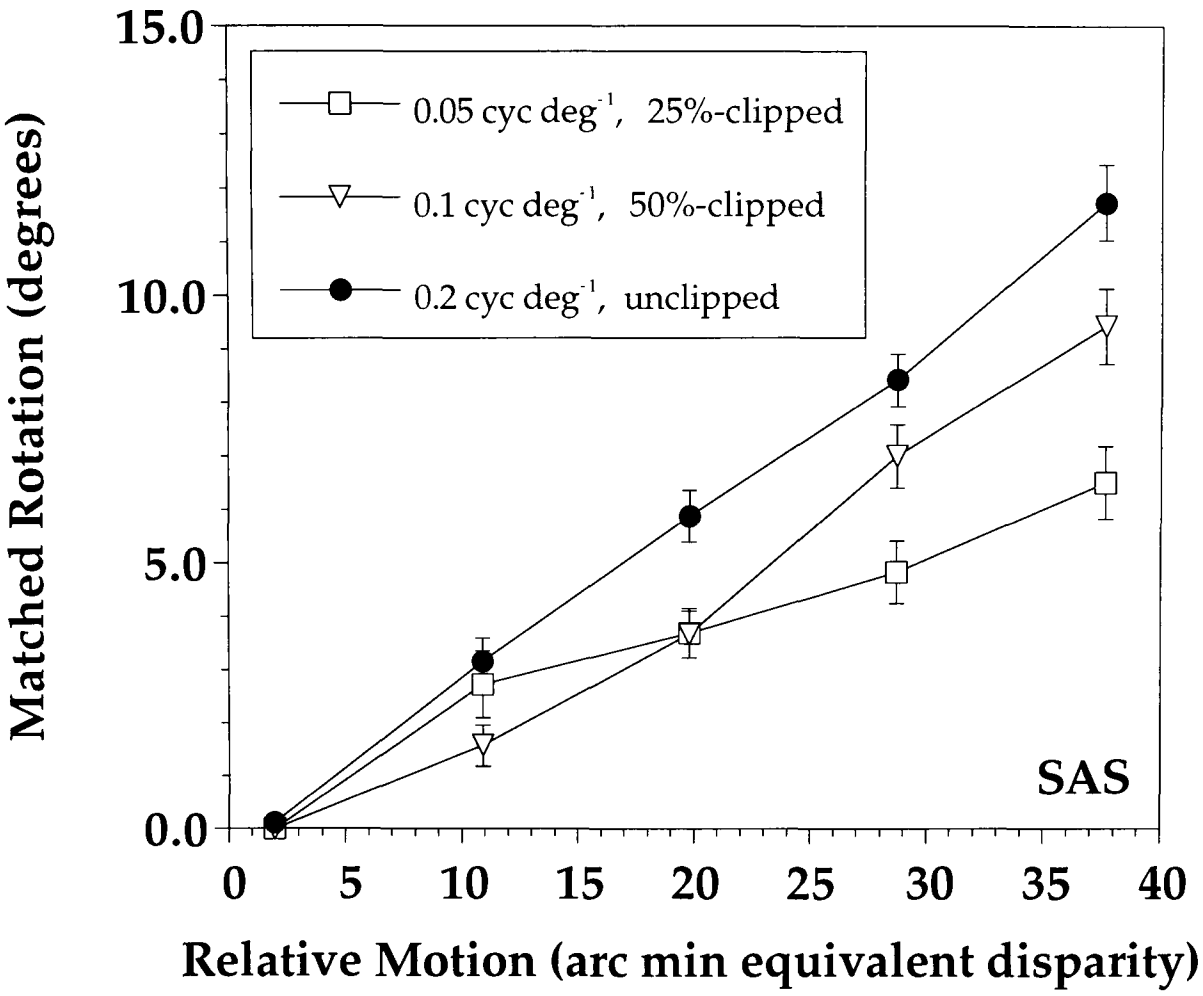


Fig. 4.26: Matched Rotation as a function of Relative Motion for surfaces containing a MMG identical to that of an unclipped 0.2 cpd surface, for Subject SAS.

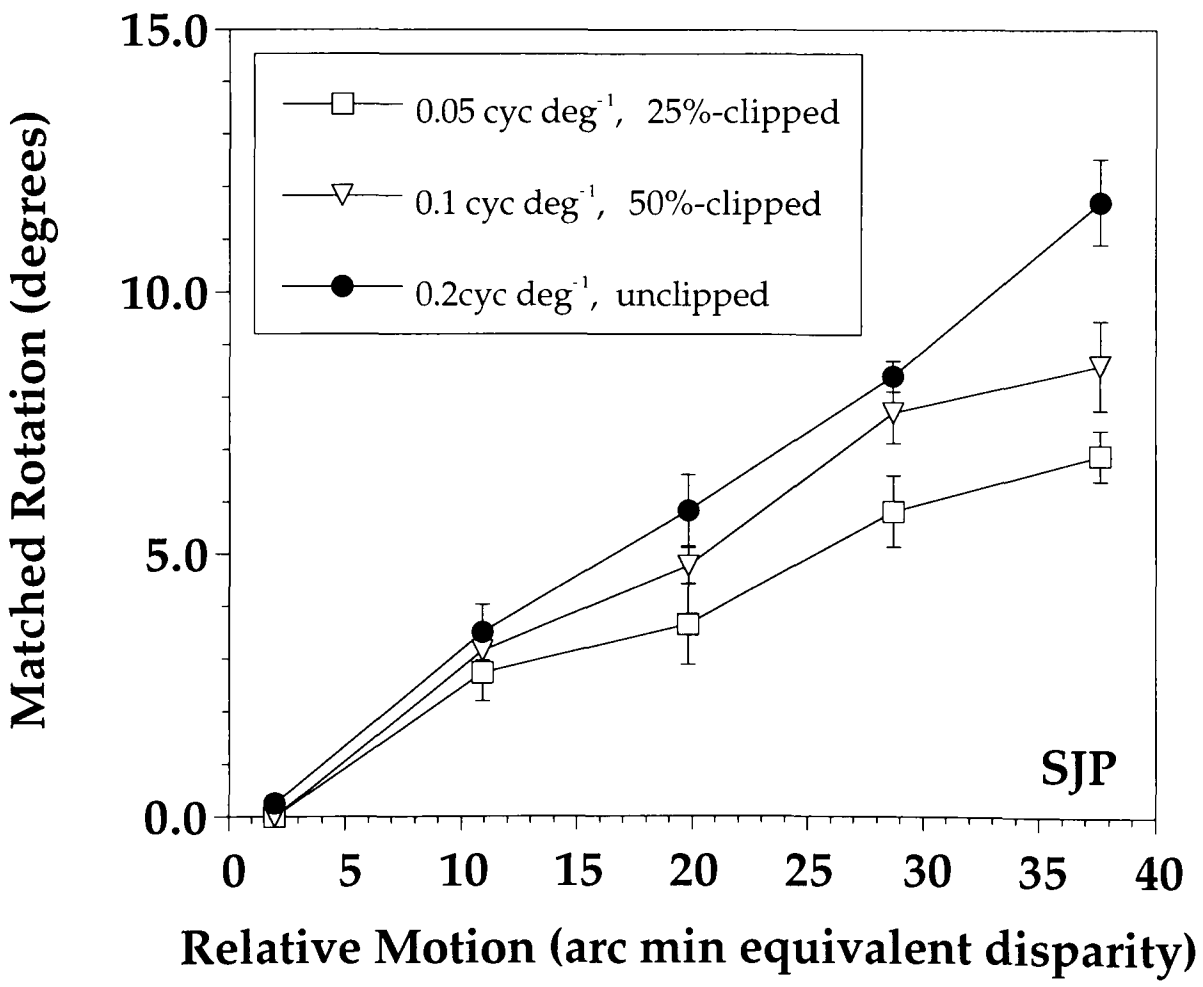


Fig. 4.27: Matched Rotation as a function of Relative Motion for surfaces containing a MMG identical to that of an unclipped 0.2 cpd surface, for Subject SJP

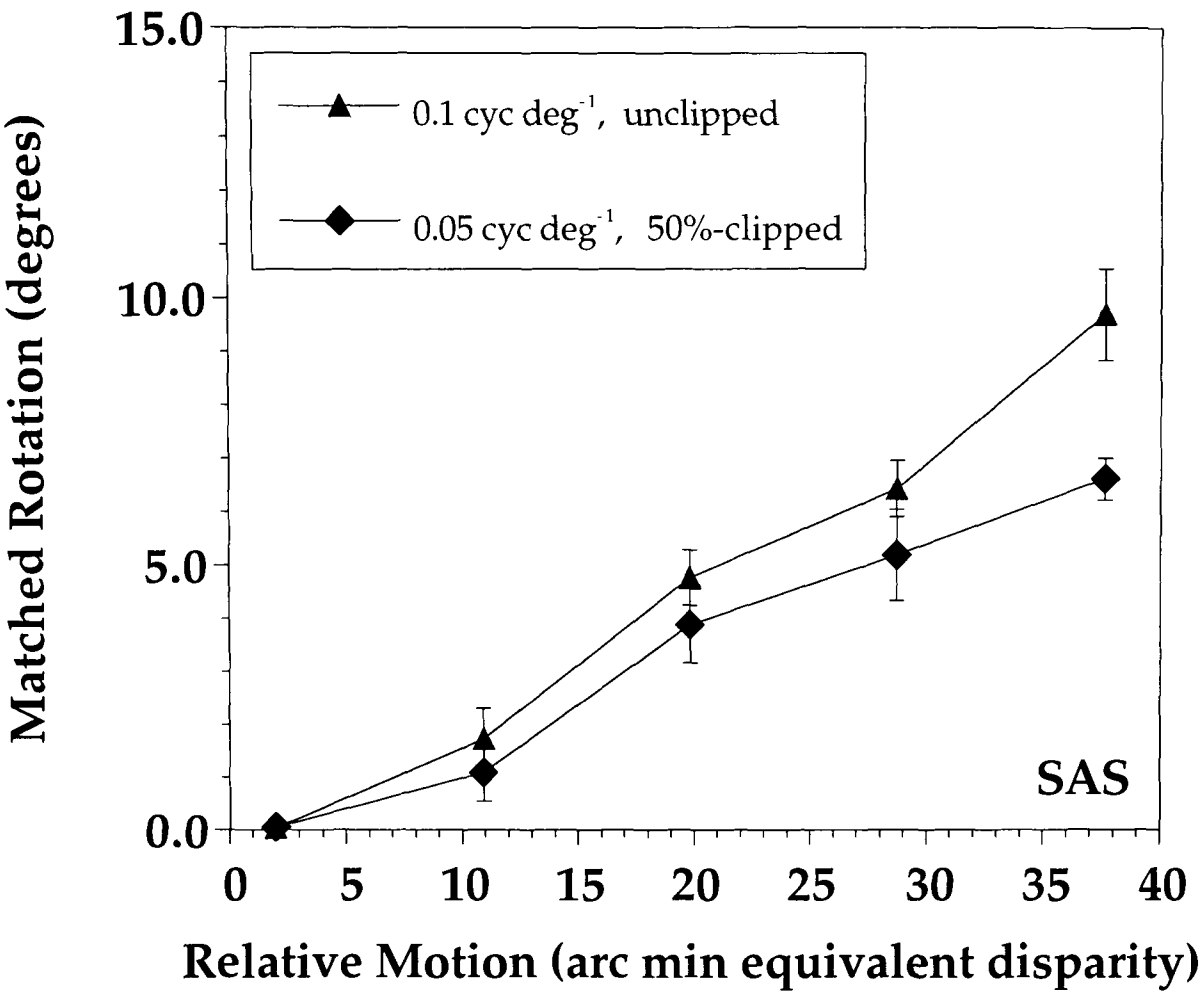


Fig. 4.28: Matched Rotation as a function of Relative Motion for surfaces containing a MMG identical to that of an unclipped 0.1 cpd surface, for Subject SAS.

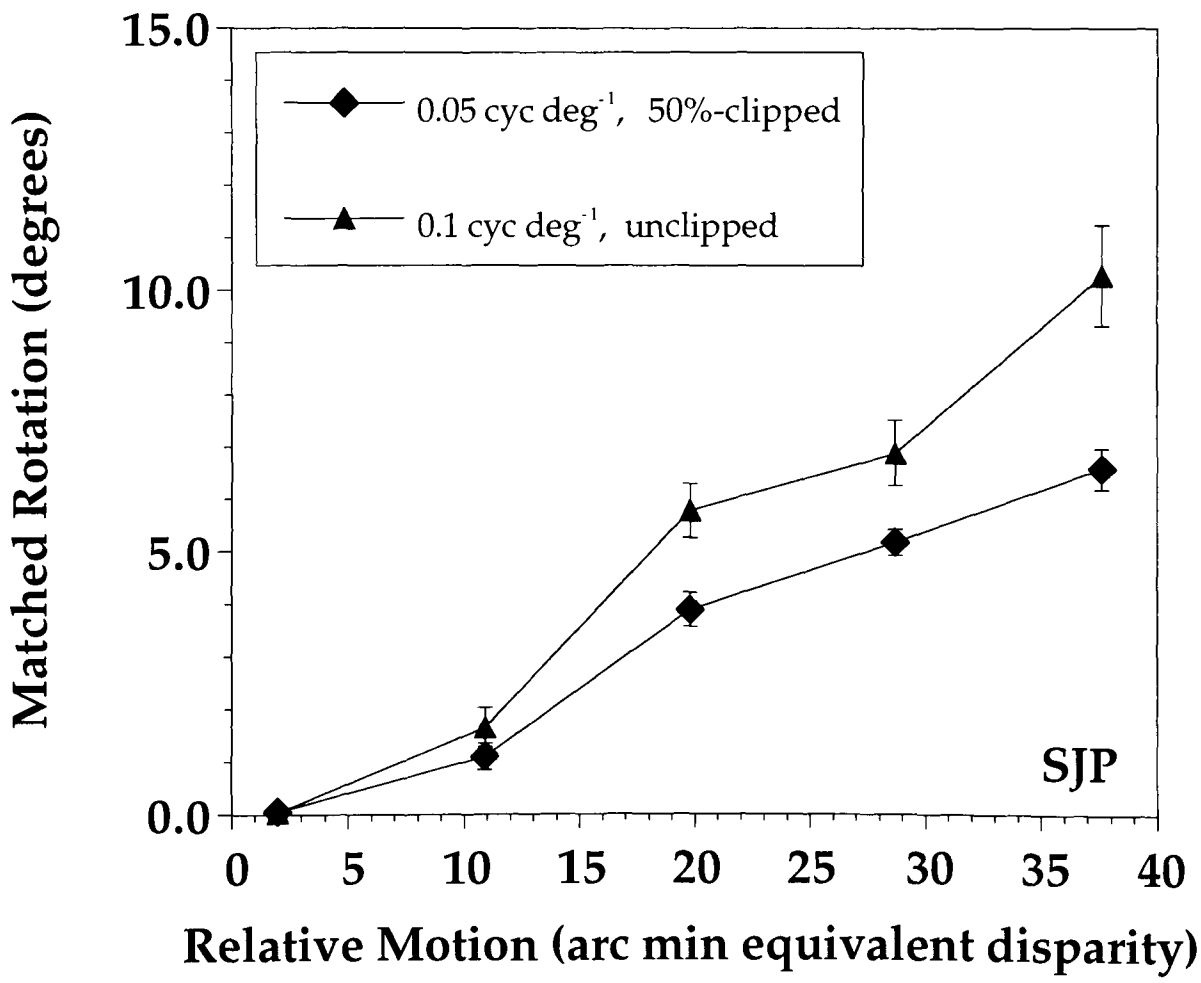


Fig. 4.29: Matched Rotation as a function of Relative Motion for surfaces containing a MMG identical to that of an unclipped 0.1 cpd surface, for Subject SJP.

4.4.5 Discussion

Increases in the MMG led to increases in rotation matches only for the stimulus of lowest corrugation frequency, 0.05 cpd. The rotation matches for stimuli of 0.1 and 0.2 cpd were generally unaffected by variations in the MMG. That the effect of MMG was evident only in the 0.05 cpd surface is not attributable to the fact that surfaces of this corrugation frequency encompassed a greater range of motion gradients than the surfaces of 0.1 and 0.2 cpd, since the effect is evident between 0.05 cpd surfaces whose MMGs differ by a factor of only 2 (e.g. rotation matches for the 50%-clipped surface are higher than those for the unclipped surface; matches for the 12.5%-clipped surface are higher than those for the 25%-clipped surface). Nor is it due to lower absolute MMGs of the 0.05 cpd surface: the 25%- and 12.5%-clipped 0.05 cpd surfaces, for example, have the same MMGs as the 50%- and 25%-clipped 0.1 cpd surfaces yet no difference was perceived between the rotation magnitudes of the latter surfaces.

This suggests that for a given amount of relative motion at the corrugation frequencies of 0.1 and 0.2 cpd, there is a maximum level of rotation which may be perceived by the visual system. When this ceiling level of rotation has been reached, the remaining 'unused' relative motion may be interpreted as non-rigidity. If this is the case, it would be expected that increases in the MMG would lead to increases in perceived non-rigidity rather than to further increases in perceived rotation. Measures of non-rigidity were not taken in the present experiment, but Williams (1993) provides data (Experiments 15 - 17 of his thesis) suggesting that non-rigidity is perceived at the expense of either rotation or depth under certain circumstances. Experiments 15 and 16 (Williams, 1993) examined the effect of discrepancies in the spatial positions of the depth corrugations specified by motion parallax and binocular disparity on perceived depth, rotation and non-rigidity. Experiment 15 employed surfaces in which the corrugation frequency specified by the parallax cue took values of 0.07 - 0.28 cpd whilst the frequency specified by the disparity cue was always 0.14 cpd. It was found that perceived rotation was greatest when both cues specified a corrugation frequency of 0.14 cpd, and that perceived rotation decreased in magnitude as the discrepancy between the frequencies specified by each cue increased. Conversely, perceived non-rigidity was zero when both cues specified a frequency of 0.14 cpd, and rose

as the discrepancy between the frequencies increased.

In Experiment 16, the corrugation frequencies specified by the parallax and disparity cues were identical (0.14 cpd), but a phase angle was introduced between the two depth profiles. This took one of nine possible values between 0° (cues in phase) and 180° (cues out of phase). It was found that perceived rotation was least at a phase angle of 90°. Perceived non-rigidity was maximal at this angle. Williams (1993) proposed that, ".. at least in the case of strongly non-rigid surfaces, there is a trade-off between rotation and non-rigidity..".

Experiment 17 used surfaces containing parallax amplitudes of 0 - 24.6 arc min equivalent disparity. There were three viewing conditions: monocular, binocular at zero disparity, and binocular at 40 arc min disparity. Observers were required to decide whether the surface 'contained independent relative motion between rows of dots'. This criterion was used as the first symptom of non-rigidity. The 50% response point (where 'rigid' and 'non-rigid' percepts were equally likely) was calculated to give the 'motion detection threshold'. Perceived depth and rotation were also measured.

This threshold was similar for the monocular and the 40 arc min binocular disparity conditions (11 - 12 arc min equivalent disparity), but was lower (6 - 10 arc min) in the zero disparity condition. Perceived rotation did not have a clear relationship with the motion threshold, but for the monocular and zero disparity conditions, perceived depth fell off as the parallax cue was increased above the motion threshold. This suggests that beyond the motion threshold, as the magnitude of the relative motion increases ".. so does the extent to which it is attributed to relative motion between rows of dots rather than to depth" (Williams, 1993, p. 324). Perceived rotation remained relatively unaffected by changes in relative motion amplitude and this could have been caused by observers' tendency to interpret relative motion as non-rigidity rather than as rotation.

The principal message of Williams' data seems to be that perceived depth and rotation may be lower than expected in some situations, and that this is generally due to the perception of non-rigidity in the surface. It

is thus possible that, for the 0.1 and 0.2 cpd surfaces of the current experiment, non-rigidity was perceived at the expense of rotation.

It was also found that stimuli of different corrugation frequencies but equal MMGs and absolute relative motion amplitudes did not result in equal rotation matches. For a given amount of relative motion, matches increased with increasing corrugation frequency. These results contrast with those of Experiment 1b, which demonstrated that there was a range of combinations of corrugation frequency and MMG within which equal MMGs resulted in approximately equal Transition Points. That experiment showed for example that a 50%-clipped 0.05 cpd surface and an unclipped 0.1 cpd surface had equal Transition Points, as did 25%-clipped 0.1 cpd and 50%-clipped 0.2 cpd surfaces, and 25%-clipped 0.05 cpd, 50%-clipped 0.1 cpd and unclipped 0.2 cpd surfaces. It was concluded that within this range the Transition Point was largely dependent on the MMG.

The results of the present experiment illustrate that for relative motion values greater than the Transition Point, perceived rotation is *not* dependent on the MMG. If this were true then the functions within each of Figs. 4.22 - 4.29 would be identical, since all surfaces within each figure contain identical MMGs.

Experiment 2 demonstrated that the Transition Point functions of square and ramp waveforms were best characterised by taking into account both the MMGs and the non-maximal motion gradients. Clipped stimuli contain both gradient discontinuities and areas of uniform motion, and thus possess characteristics of square waveforms. The results of Experiment 1b showed that square waveforms have higher Transition Points than sine waveforms of the same corrugation frequency. This implies that for a given amount of relative motion less rotation is perceived in square waves than in sine waves. This suggestion is supported by the finding that observers do not generally perceive rotation in square wave stimuli, but rather perceive a 'shearing motion' (Ono and Steinbach, 1990). Since clipped stimuli resemble both sine and square waveforms, it would be predicted that they would give rise to both shearing and rotation percepts, with the relative weightings of these percepts being dependent on the similarity of the waveform to either the sine or the square waveform. For a given motion gradient, the lower the

corrugation frequency the greater the degree of peak-clipping and thus the greater the area of uniform motion. For example, an unclipped stimulus of frequency 0.2 cpd contains no regions of uniform motion; a 0.1 cpd 50%-clipped stimulus contains regions of uniform motion which cover over half its total 2-D area; and a 0.05 cpd 25%-clipped stimulus consists almost entirely of regions of uniform motion.

The decreased amount of perceived rotation in low frequency stimuli may arise through the relatively high weighting given to the shearing percept in these surfaces, since their large areas of uniform motion result in their close resemblance to a square waveform. If this is the case, there is consequently less relative motion available for interpretation as perceived rotation (assuming that there is some form of trade-off between the magnitudes of the possible percepts).

Were this to be true, then the reason for the increase in perceived rotation with increasing corrugation frequency and concurrent decreasing clipped-ness might be that the observer attributes increasing amounts of relative motion to non-rigidity as the degree of clipped-ness increases. Assuming that relative motion cannot be interpreted as both non-rigidity *and* rotation, the former may be perceived at the expense of the latter. As discussed above, Williams (1993) has shown that this trade-off does indeed occur in some situations.

4.5 Experiment 8: Matched Rotation as a function of Stimulus Form

4.5.1 Introduction

Experiments 7a and 7b demonstrated that the MMG plays a role in the magnitude of perceived rotation, albeit a smaller one than in the determination of Transition Points. The present experiment examined whether perceived rotation is influenced by the sign of the central corrugation.

4.5.2 Stimuli

Stimuli contained a corrugation frequency of 0.05, 0.1 or 0.2 cpd. There were two types of stimulus at each corrugation frequency: one contained a central peak, the other contained a central trough. Five values of relative motion were used.

4.5.3 Procedure

Observers' eye movements were unrestricted, but they were instructed to attend to the central corrugation when making their rotation match. The procedure was otherwise identical to that of Experiment 7a, except that subjects completed four sessions of one block per corrugation frequency per session.

4.5.4 Results

Figs. 4.30 - 4.38 show Matched Rotation as a function of relative motion amplitude, for the three corrugation frequencies. The filled symbols represent the matches obtained with stimuli containing a central peak. The open symbols represent the matches obtained with stimuli containing a central trough.

Data were pooled across subjects and were entered into three separate two-way ANOVAs, one for each corrugation frequency. The factors in each ANOVA were central corrugation (peak/trough) and relative motion amplitude (4, 12, 20, 28, 36 arc min equivalent disparity).

For the 0.05 cpd surface, perceived rotation was generally greater when the surface contained a central peak than when it contained a central trough. This difference was shown to be significant ($F[1, 2] = 57.48$, $p<0.05$). The main effect of relative motion amplitude was significant ($F[4, 8] = 25.04$, $p<0.0005$). The interaction between central corrugation and relative motion amplitude was significant ($F[4, 8] = 2.62$, $p<0.05$).

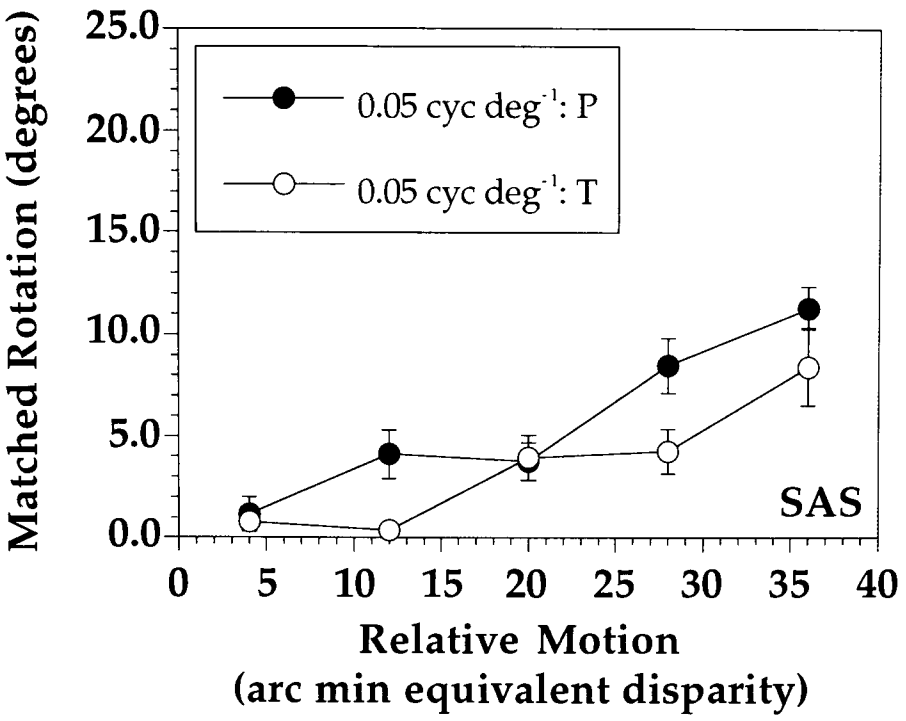
For the 0.1 cpd surface, perceived rotation was generally greater when the surface contained a central trough than when it contained a central peak. This difference was significant ($F[1, 2] = 39.77$, $p<0.05$). The main

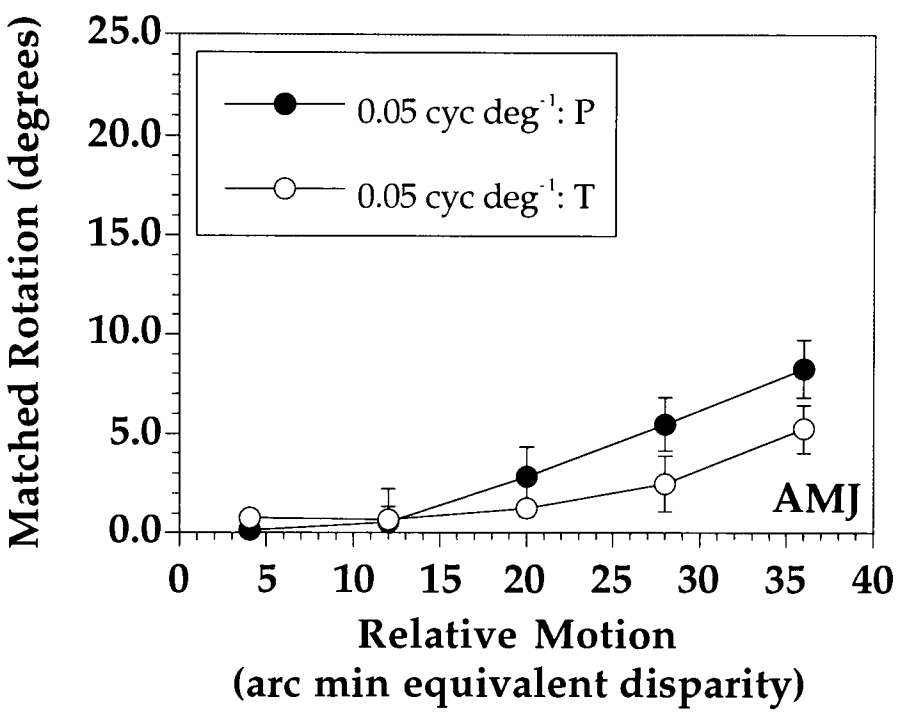
effect of relative motion amplitude was significant ($F[4, 8] = 37.06$, $p<0.0005$). The interaction between central corrugation and relative motion amplitude was significant ($F[4, 8] = 3.88$, $p<0.05$).

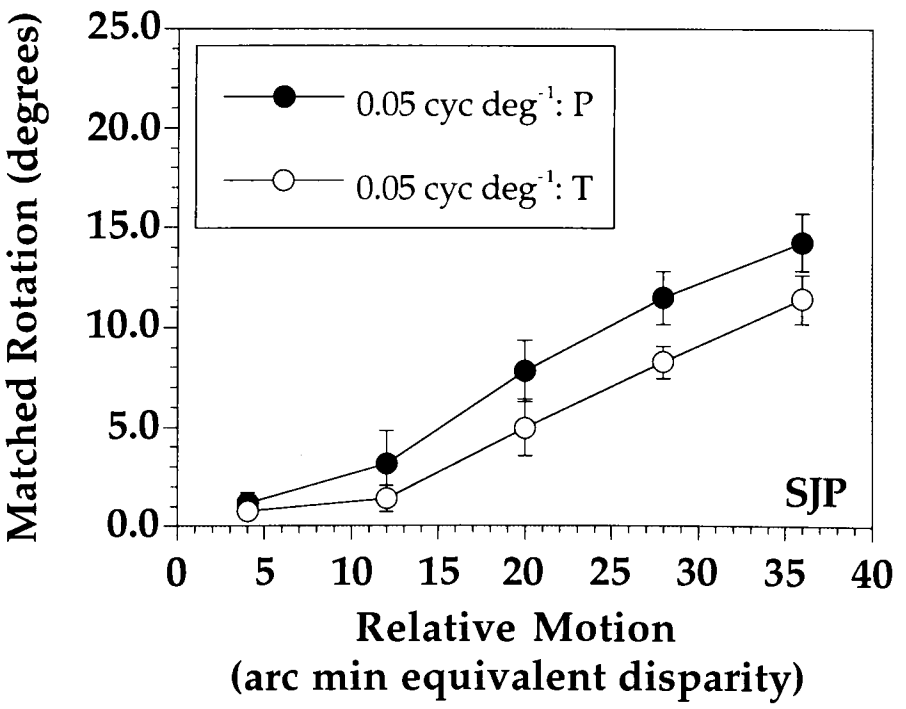
For the 0.2 cpd surface, the main effect of the sign of the central corrugation was non-significant ($F[1, 2] = 3.71$, n.s.). The main effect of relative motion amplitude was significant ($F[4, 8] = 67.53$, $p<0.0005$). The interaction between central corrugation and relative motion amplitude was non-significant ($F[4, 8] = 3.74$, n.s.).

4.5.5 Discussion

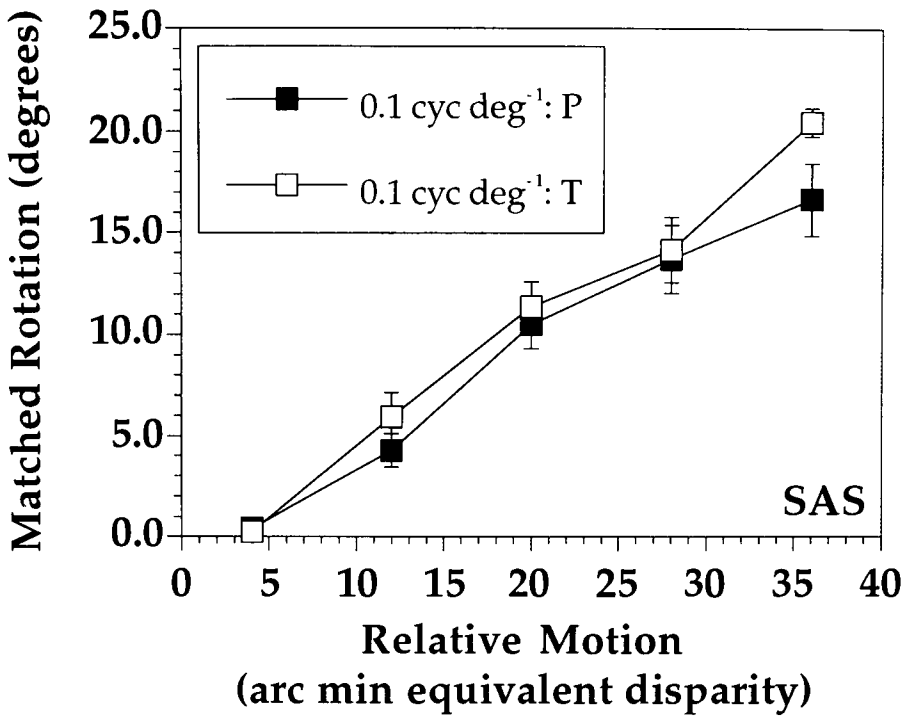
The finding that more rotation was perceived in 0.05 cpd surfaces containing a central peak than in those containing a central trough suggests that stimulus form affects the magnitude of perceived rotation under certain circumstances. That form does not have an identical effect upon the magnitude of perceived rotation across corrugation frequencies is evident from the data of the 0.1 and 0.2 cpd surfaces, the pattern of which is qualitatively and quantitatively different to that of the 0.05 cpd data. The qualitative difference (namely that 0.05 cpd surfaces containing central peaks were perceived to rotate more than those containing central troughs whilst the reverse was true for 0.1 cpd surfaces) suggests that it is not the sign of the central corrugation in itself which determines the relative magnitudes of perceived rotation.

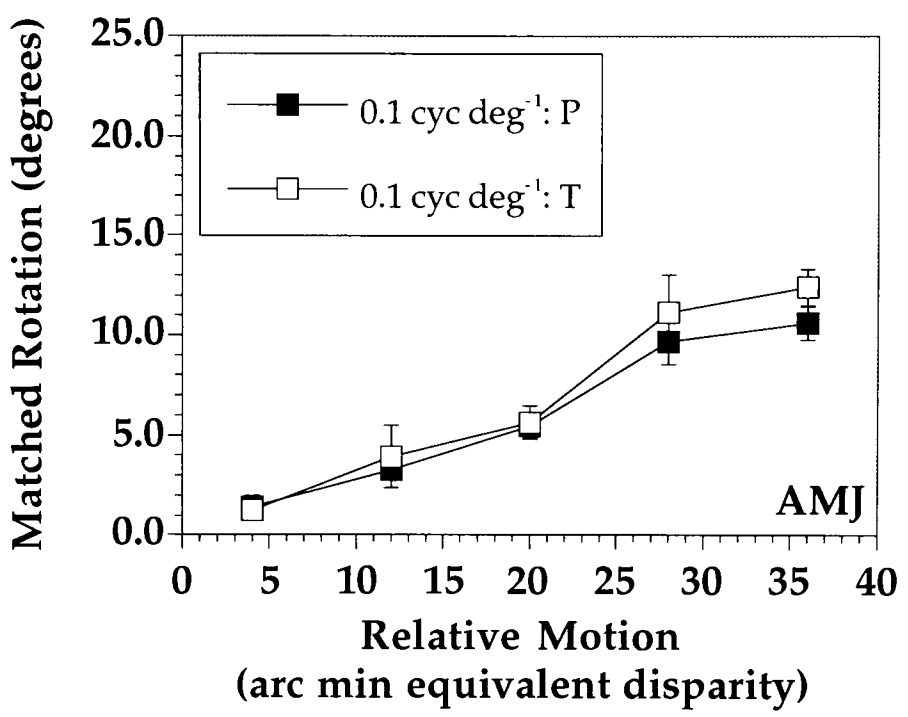


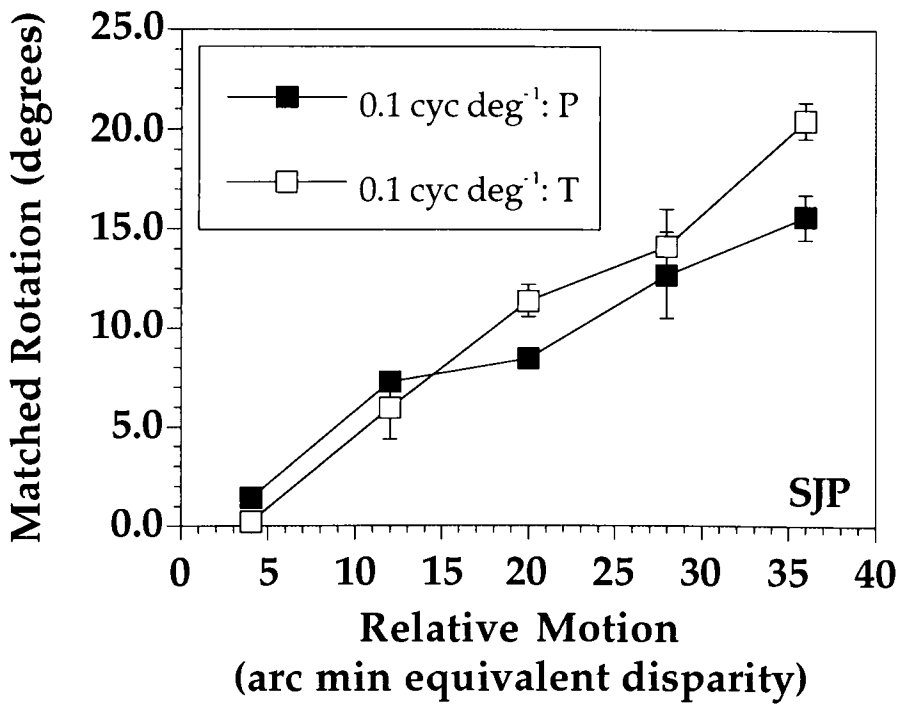




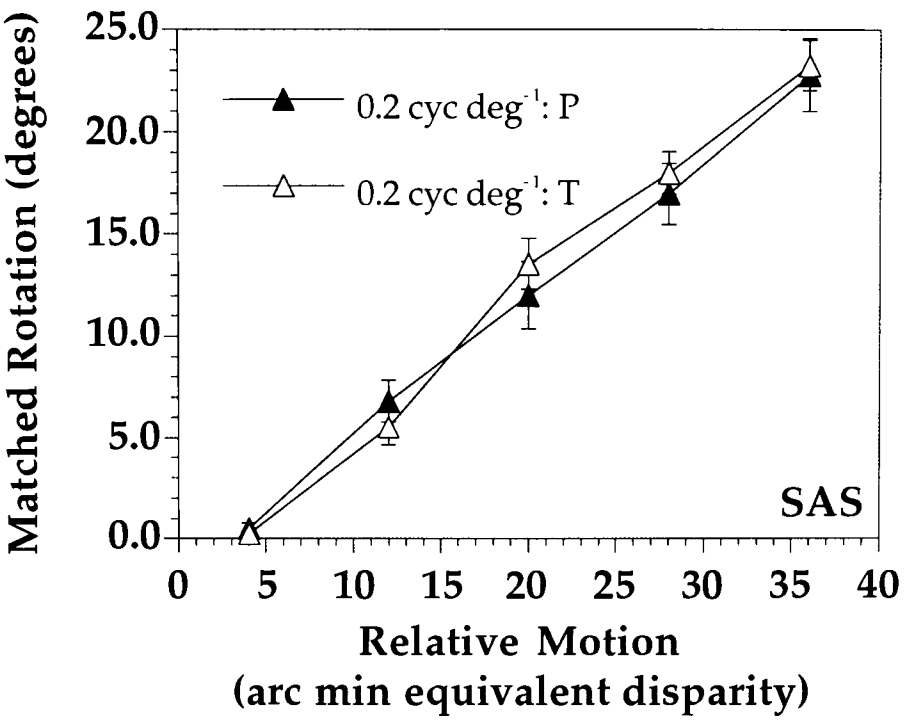
Figs. 4.30 - 4.32: Matched Rotation as a function of Relative Motion for a 0.05 cpd surface, for Subjects SAS, AMJ and SJP. Filled symbols represent surfaces containing a central peak; open symbols represent surfaces containing a central trough.

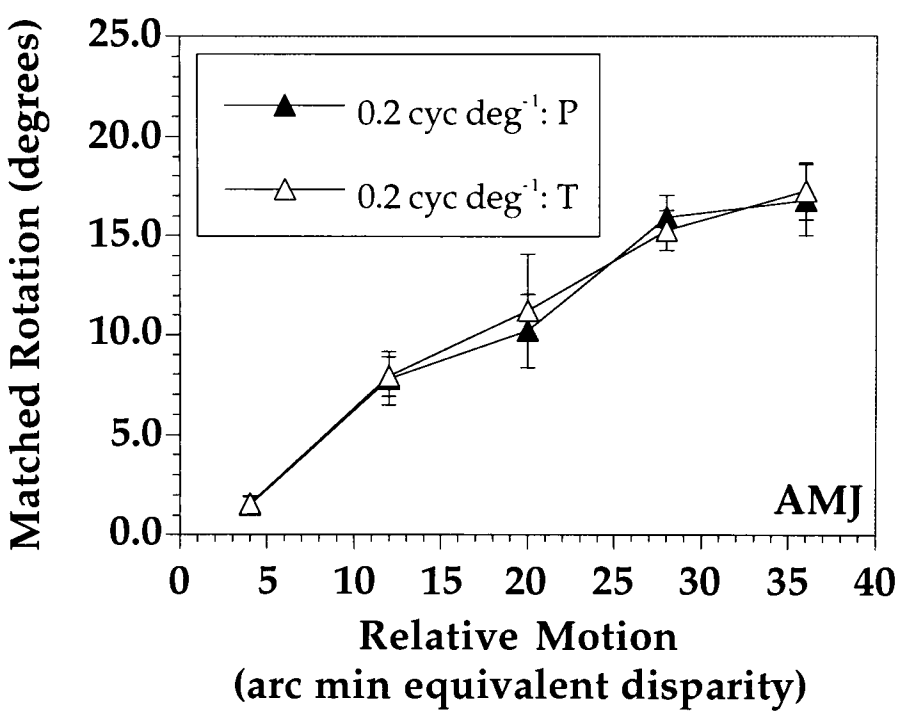


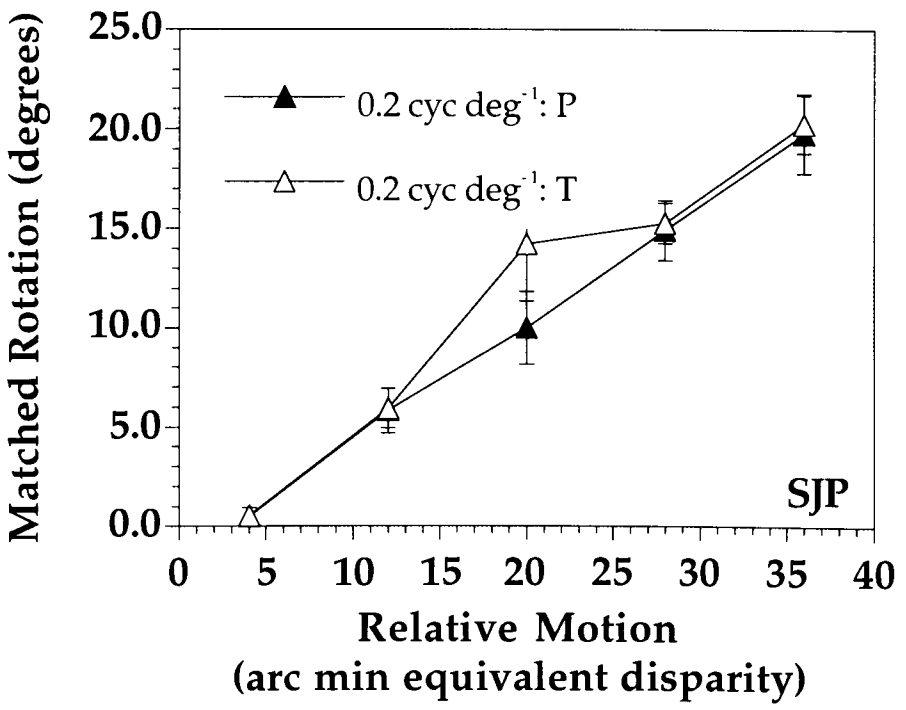




Figs. 4.33 - 4.35: Matched Rotation as a function of Relative Motion for a 0.1 cpd surface, for Subjects SAS, AMJ and SJP. Filled symbols represent surfaces containing a central peak; open symbols represent surfaces containing a central trough.







Figs. 4.36 - 4.38: Matched Rotation as a function of Relative Motion for a 0.2 cpd surface, for Subjects SAS, AMJ and SJP. Filled symbols represent surfaces containing a central peak; open symbols represent surfaces containing a central trough.

As specified in Section 4.5.3, observers' eye movements were unrestricted, but they were instructed to attend to the central corrugation whilst making the rotation judgement. The distance between that attended-to central corrugation and the adjacent corrugations of opposite sign obviously decreases with increasing corrugation frequency. Hence it might be expected that the influence of the non-central corrugations on perceived rotation would increase with their proximity to the central corrugation. If this were true then it would explain the increasing similarity of the magnitudes of perceived rotation in stimuli containing central corrugations of opposite sign with increasing corrugation frequency — recall that there was a significant main effect of the sign of the central corrugation in 0.1 cpd, but not 0.2 cpd, surfaces. However, it does not explain the qualitative difference between the data of the 0.05 cpd surface and the data of the 0.1 and 0.2 cpd surfaces.

Braunstein and Andersen (1981) presented observers with a translating dihedral structure composed of two planes which met at a horizontal line in the centre. This stimulus can be thought of as a single cycle of a 0.03 cpd triangle wave (the visible part of the stimulus was a disc of diameter 1.2 m and viewing distance was 1.8 m). Subjects were asked whether the centre appeared to be the nearest or the furthest part of the display, and also to classify the rotation perceived as "no rotation", "some rotation" or "mostly rotation". No effect of the direction of motion gradient was found, a result which is inconsistent with that of the current experiment. Several differences obviously exist between the Braunstein and Andersen (1981) study and the current one, however: the nature of the waveform, the viewing distance, the fact that the parallax was not linked to head motion in the former study, and fixation instructions at the time of rotation judgement. Perhaps more importantly, the method used to quantify the magnitude of perceived rotation in the former experiment was rather crude. Since the differences between the perceived rotation of surfaces with opposite velocity gradients found in the current experiment were relatively small, it is possible that such differences may not have been detectable using the rotation classification procedure of Braunstein and Andersen (1981).

4.6 Experiment 9: Matched Rotation as a function of Stimulus Size

4.6.1 Introduction

A number of studies has demonstrated the importance of vertical perspective information in the determination of depth order and rotation direction (e.g. Power, 1967; Braunstein and Payne, 1968; Börjesson, 1971; Hershberger, Stewart and Laughlin, 1976; Braunstein, 1977b; Rogers and Rogers, 1992). Many studies have illustrated that perspective is also an important factor in quantitative estimations of slant (Clark, Smith and Rabe, 1955; Gibson and Gibson, 1957; Gillam, 1968; Youngs, 1976; Gillam and Ryan, 1992; Eby and Braunstein, 1993; Ryan and Gillam, 1994). These studies are considered fully in Chapter 6 (see also Section 3.8.4).

Experiment 3 showed that the Transition Point increased with increasing stimulus size. It was hypothesised that the reason for the lower Transition Points of the smaller stimuli was the smaller absolute size of the vertical perspective cues which signalled that the surface remained in the frontal plane. The perception of rotation was thus facilitated as stimulus size decreased, because the percept involved a smaller conflict with the perspective cue. The present experiment examined the effect of stimulus size on the magnitude of perceived rotation for surfaces with relative motion amplitudes greater than the Transition Point. It is predicted that for a given amount of relative motion, more rotation will be perceived in smaller than larger stimuli.

4.6.2 Stimuli

Stimuli were of sizes 2.5, 5, 10 and 20°. The corrugation frequency was 0.4 cpd. Seven values of relative motion were used.

4.6.3 Procedure

This was identical to that of Experiment 7a.

4.6.4 Results

Figs. 4.39 - 4.41 show rotation matches for each of the four stimulus sizes. Perceived rotation increased with increasing relative motion for all sizes tested. The gradient of the function increased with decreasing size, as is evident from Figs. 4.42 - 4.44, which show the lines of best fit for these data, and their equations.

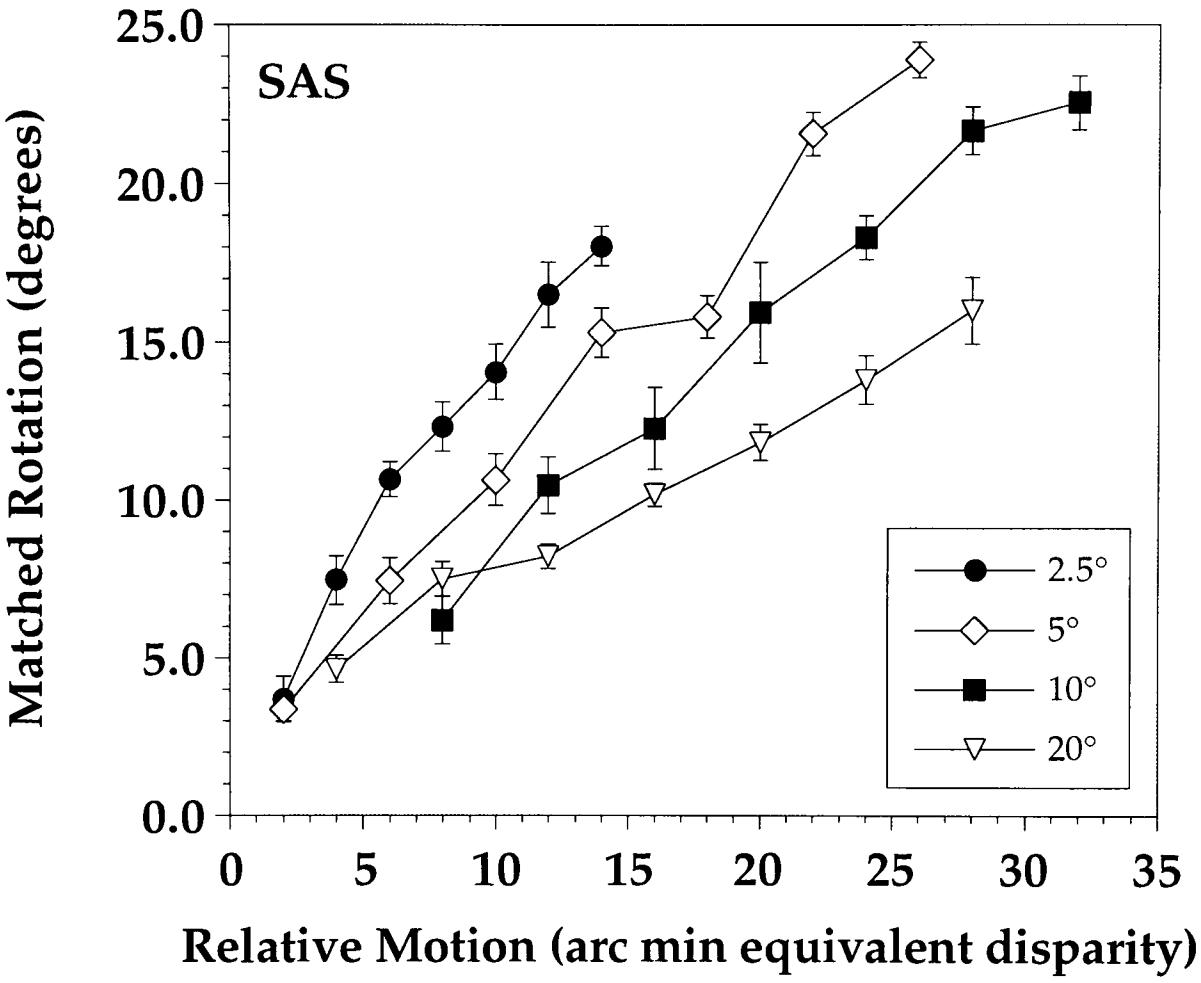


Fig. 4.39: Matched Rotation as a function of Relative Motion for four Stimulus Sizes, for Subject SAS.

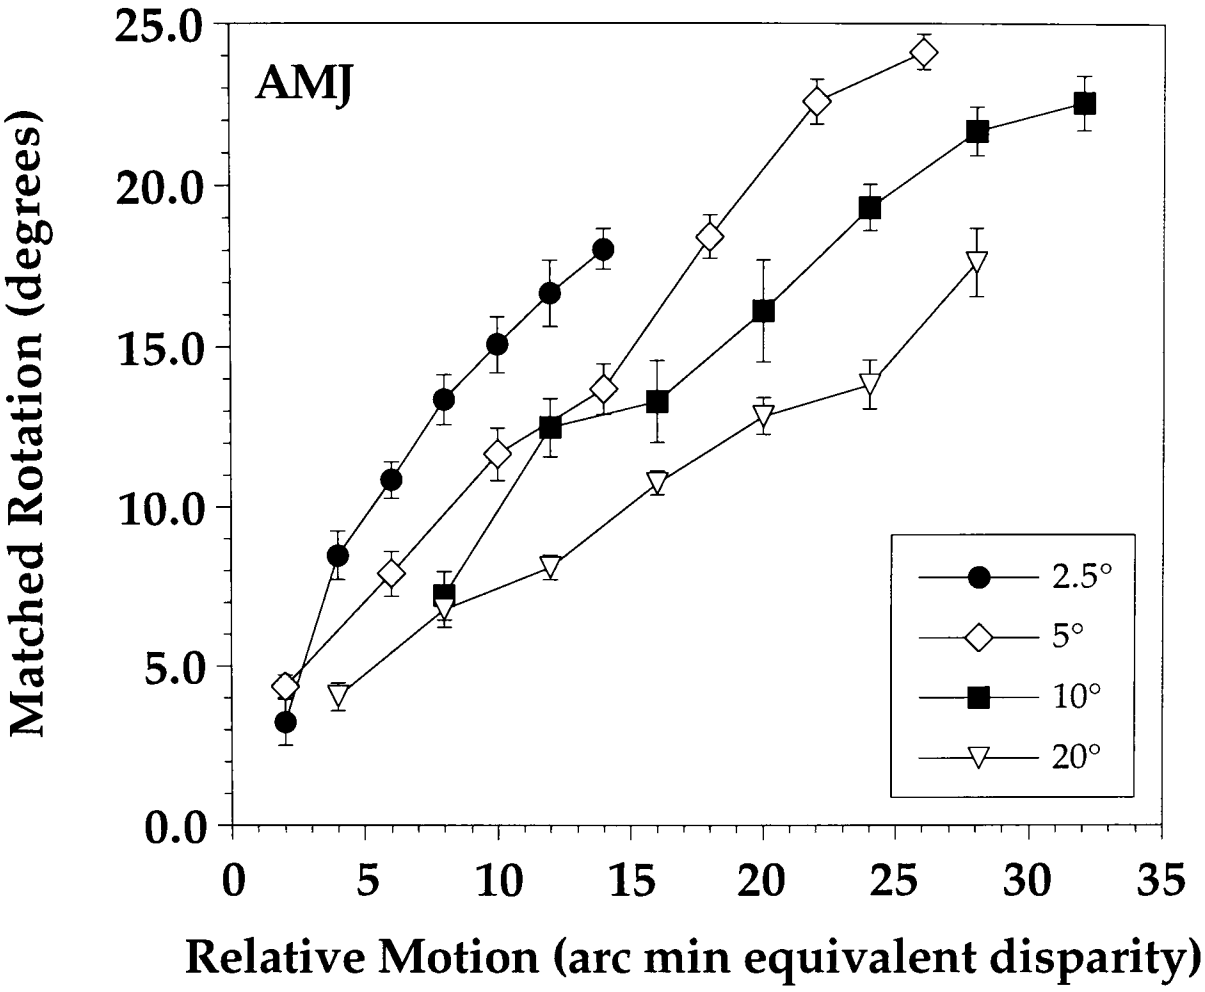


Fig. 4.40: Matched Rotation as a function of Relative Motion for four Stimulus Sizes, for Subject AMJ.

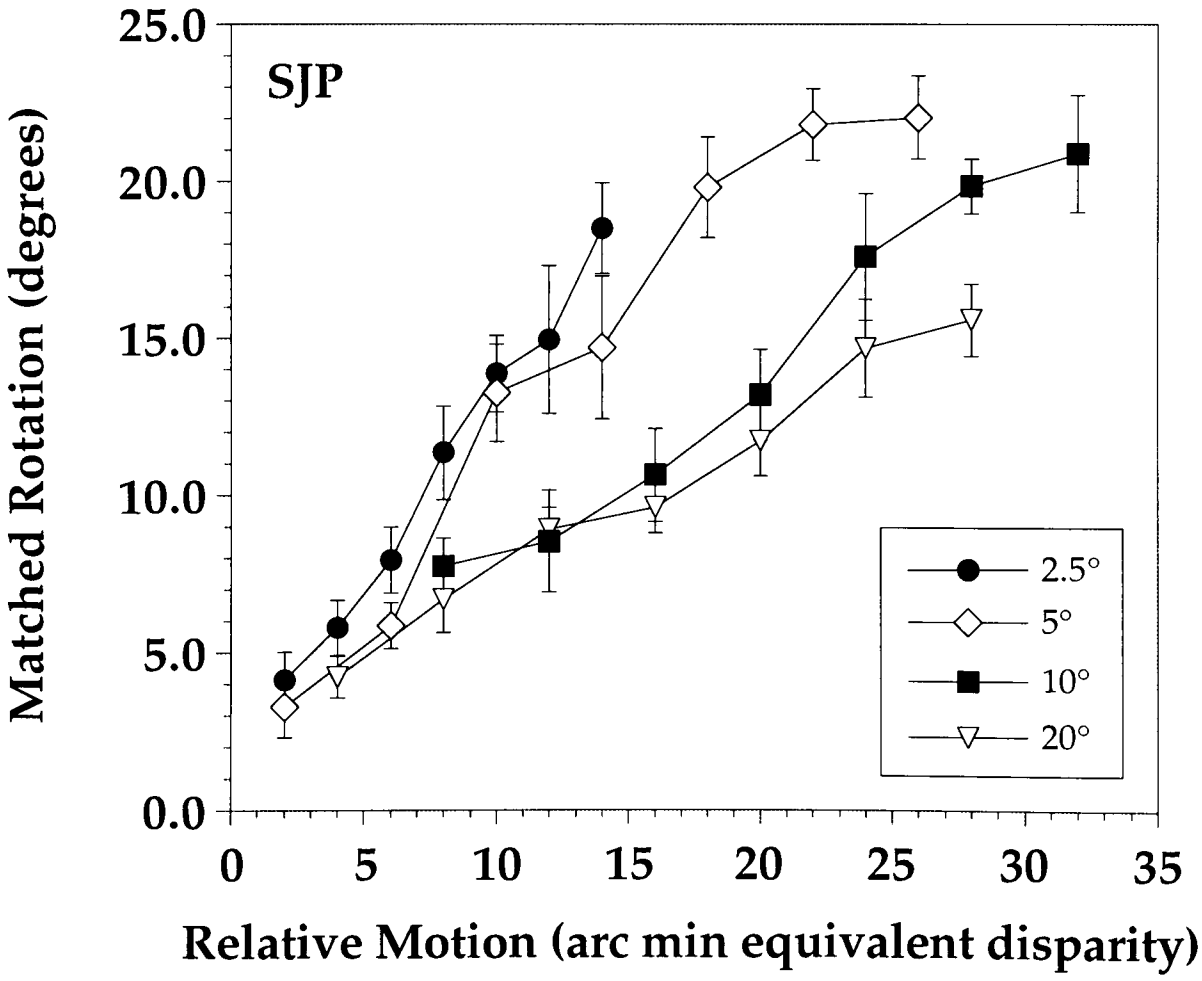
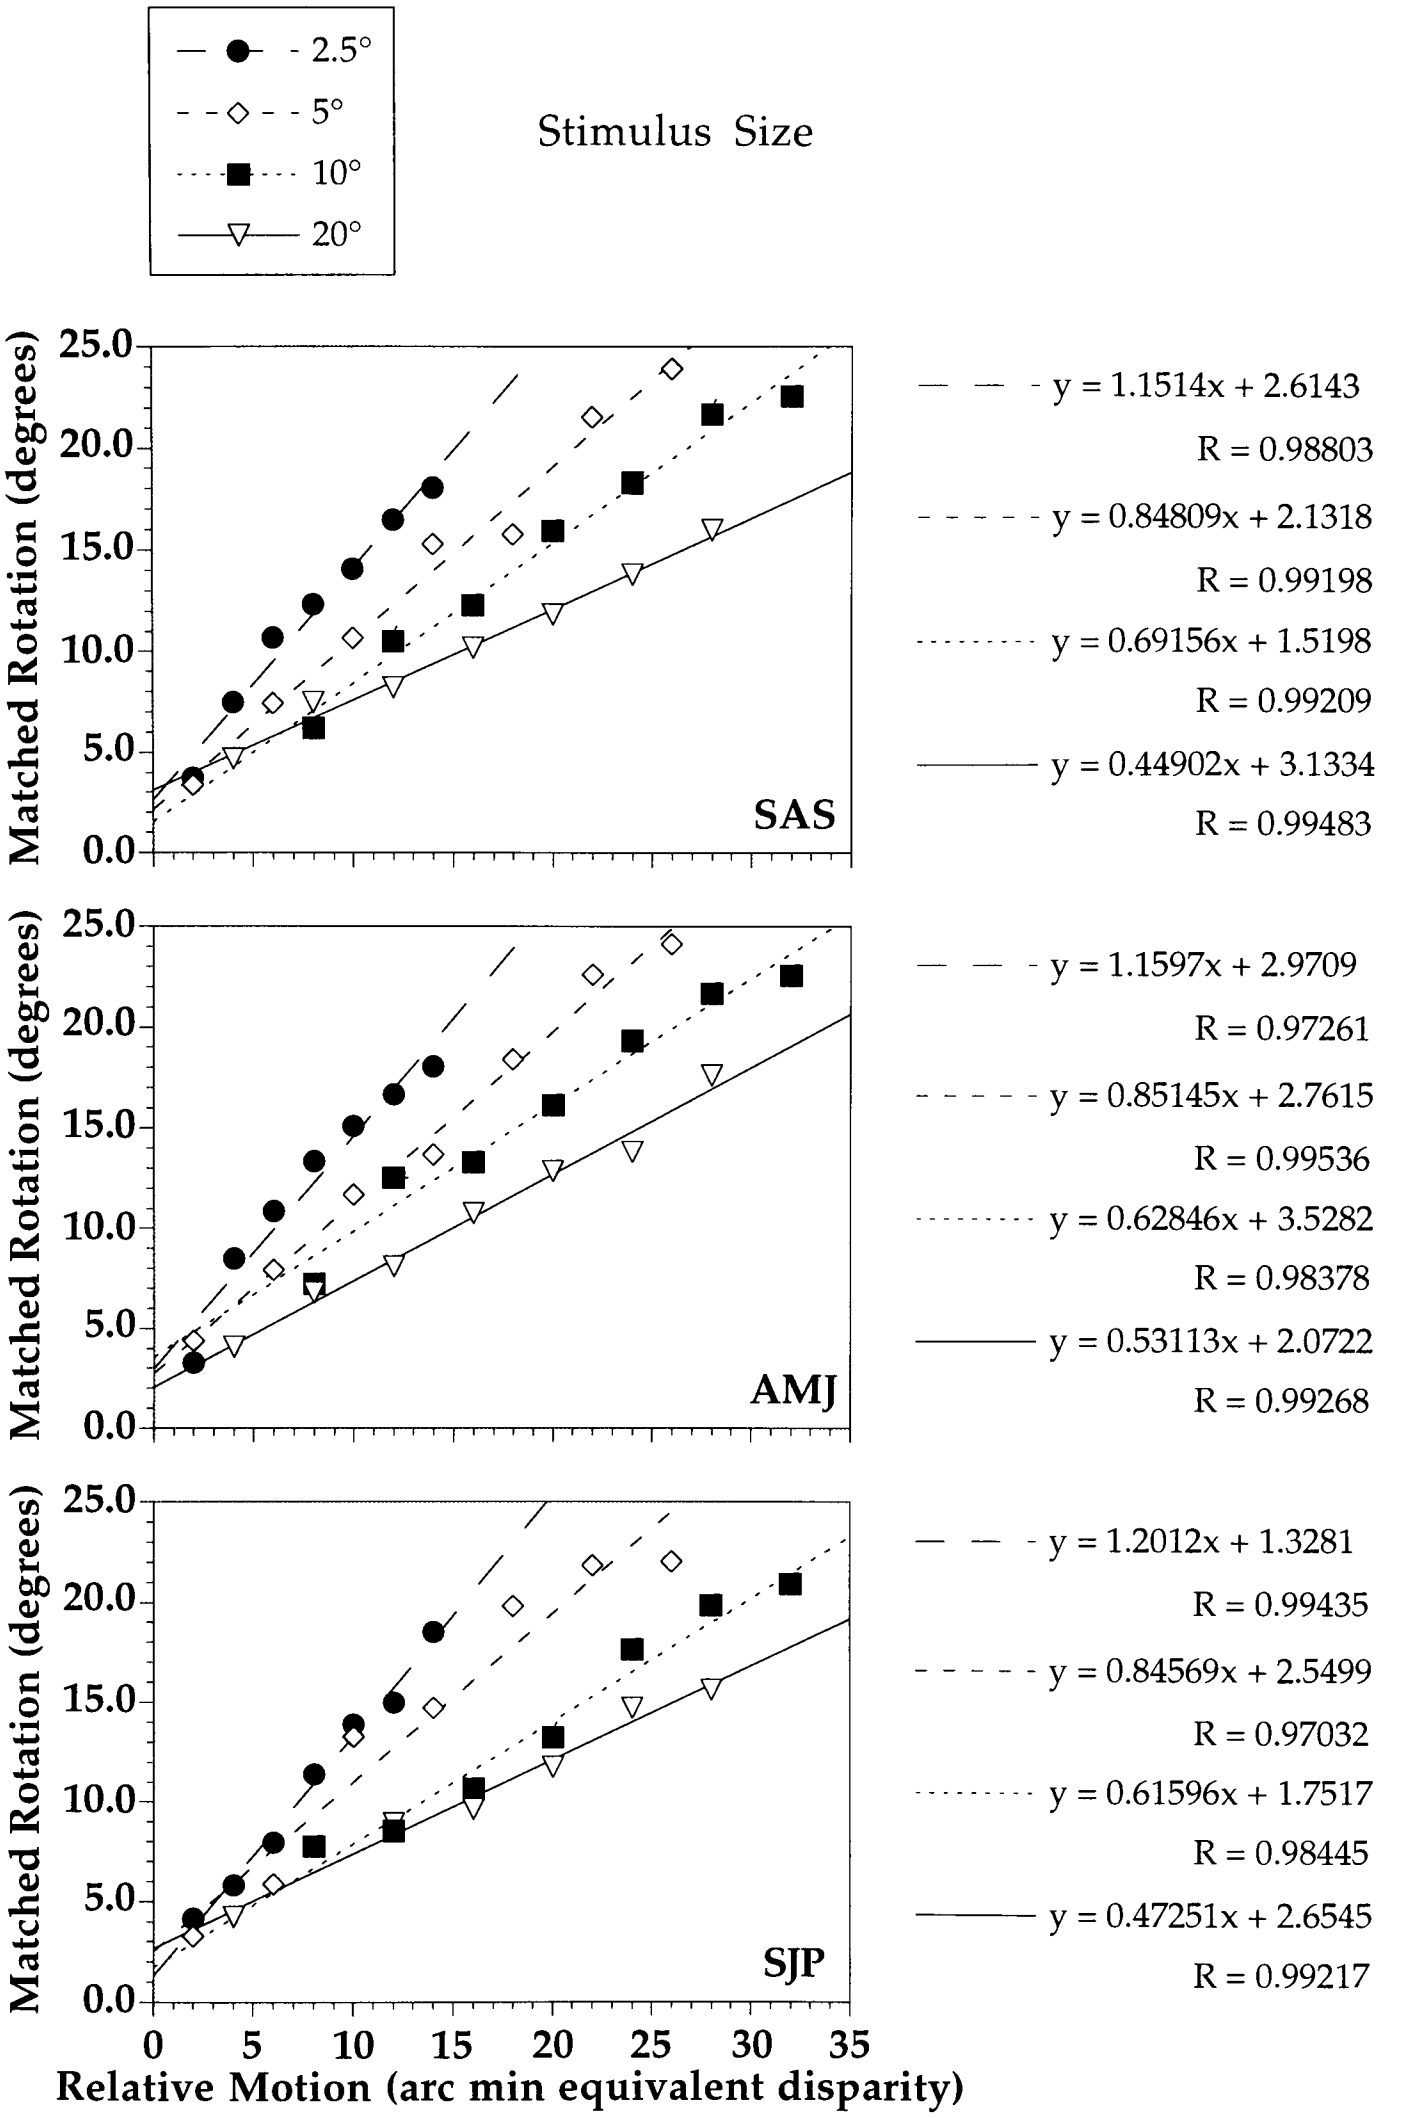


Fig. 4.41: Matched Rotation as a function of Relative Motion for four Stimulus Sizes, for Subject SJP.



Figs. 4.42 - 4.44: Matched Rotation as a function of Relative Motion for the four stimulus sizes, with the lines of best fit and their equations for Subjects SAS, AMJ and SJP.

4.6.5 Discussion

For a surface at a given viewing distance, the absolute amount of vertical perspective change for a given rotation angle increases with increasing stimulus size. When the vertical perspective information specifies zero rotation (as in the current experiment), any perception of rotation involves a conflict with the vertical perspective information. On this basis, it would be expected that a decreasing amount of rotation would be perceived as stimulus size increases. This is exactly the pattern of results obtained: for a given amount of relative motion, the magnitude of perceived rotation is greater in small than in large stimuli.

Experiment 3 showed that for each doubling of stimulus size the Transition Point increased by a factor of approximately 0.2. The corresponding factor of the current experiment, averaged across subjects, was 0.35. It therefore seems that the effect of stimulus size at supra-threshold rotation levels is stronger than its effect at threshold levels. This is perhaps due to an increased importance of the perspective cue in the latter situation. These results contrast with the *smaller* weighting given to the MMG at supra-threshold rotation levels than at threshold levels. It is thus evident that the relative weightings of these two cues to rotation — vertical perspective and the MMG — differ according to whether the levels of rotation are threshold or supra-threshold.

Another factor which may have played a role in the relative magnitudes of perceived rotation is the extent to which the observer perceived the stimulus as 'fixed' in the plane of the monitor. When the observer translates through 6.5 cm and the rotation of the paddle with respect to the ground is zero, the paddle and monitor undergo exactly the same amount of rotation (both with respect to the ground (0°) and with respect to the eye (6.5°)). If the stimulus is always perceived to lie in the plane of the monitor, then rotation matches of zero would be predicted. Such matches evidently were not obtained, and the degree to which the surface appeared to rotate may be an indication of the degree to which it was perceived to be 'fixed' in the plane of the screen. It is possible that the reason for the increase in perceived rotation magnitude with decreasing stimulus size is that smaller stimuli are further from the boundaries of the monitor. They may thus be perceived as being increasingly

independent of their stationary environment and, as such, are 'free' to rotate.

4.7 Experiment 10: Matched Rotation as a function of Central and Flanking Corrugation Frequencies

4.7.1 Introduction

Previous experiments of this thesis have shown that corrugation frequency plays a role in the Transition Point (Experiments 1a and 1b) and, at higher relative motion values, partly determines the magnitude of perceived rotation (Experiments 7a and 7b). The surfaces of these experiments contained a single corrugation frequency. The present experiment investigated the perceived rotation of surfaces comprising two different corrugation frequencies. The total 2-D areas occupied by each frequency were identical. One frequency occupied the central portion of the surface, the other frequency occupied the portions above and below (see Section 4.7.2).

The experiment was designed to examine the effect of the flanking frequency on the perceived rotation of surfaces with identical central frequencies, or vice versa. On the basis of previous work of this chapter it is predicted that the surface with the higher flanking (central) frequency will appear to rotate more. To anticipate the results, this prediction was generally confirmed. This experiment determined to what extent the effect is dependent upon the corrugation frequencies of the surface, and whether the size of the effect varies according to whether the identical frequencies are in the central or the flanking regions of the surface. The latter finding indicates the relative importance of the central and the outer portions of a surface in determining its perceived rotation.

4.7.2 Stimuli

Stimuli were 20° in diameter and contained two different corrugation frequencies. One frequency occupied the upper 25% and the lower 25% of the total height. The other frequency occupied the central 50% of the total height. The frequency in the latter position will henceforth be termed the

'central' frequency, and the frequency in the former position will be termed the 'flanking' frequency.

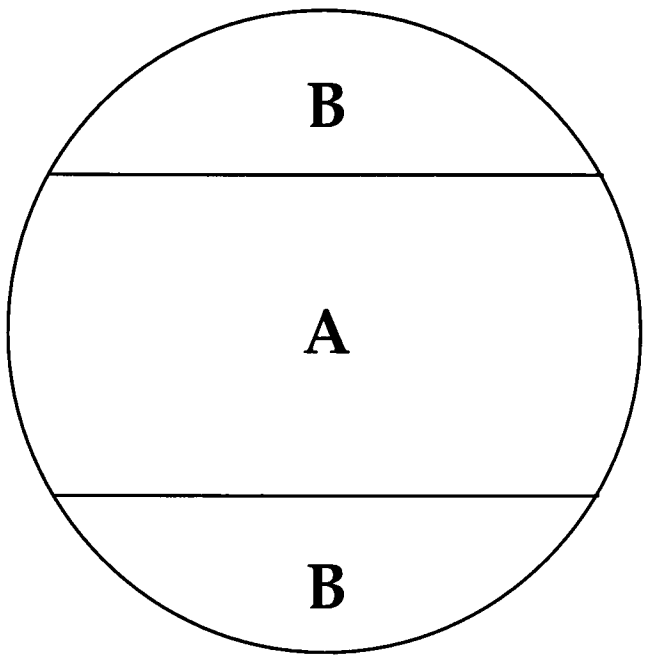


Fig. 4.45: The central 50% of the total height of the stimulus (region A) contained a different corrugation frequency to the upper 25% and lower 25% of the total height (regions B). The corrugation frequency occupying region A was termed the 'central' frequency. The frequency occupying regions B was termed the 'flanking' frequency.

The surface always contained a central peak. Various combinations of corrugation frequencies 0.05, 0.1 and 0.2 cpd were used. The combinations are shown below, with the central frequency first:-

0.05, 0.1
0.05, 0.2
0.1, 0.05
0.1, 0.2
0.2, 0.05
0.2, 0.1.

4.7.3 Procedure

This was identical to that of Experiment 7b.

4.7.4 Results

Figs. 4.46 - 4.51 show Rotation Match as a function of Relative Motion for stimuli containing identical central corrugation frequencies but different flanking frequencies. Stimuli with the higher flanking frequency

generally resulted in greater rotation matches for surfaces with a central corrugation frequency of 0.05 or 0.1 cpd. This effect was much diminished for surfaces of central corrugation frequency 0.2 cpd.

Data from each subject were entered into a three-way repeated-measures ANOVA, with central corrugation frequency (0.05, 0.1, 0.2 cpd), flanking corrugation frequency (high, low) and relative motion amplitude (4, 12, 20, 28 and 36 arc min equivalent disparity) as the factors. Trial repetition was the repeated factor.

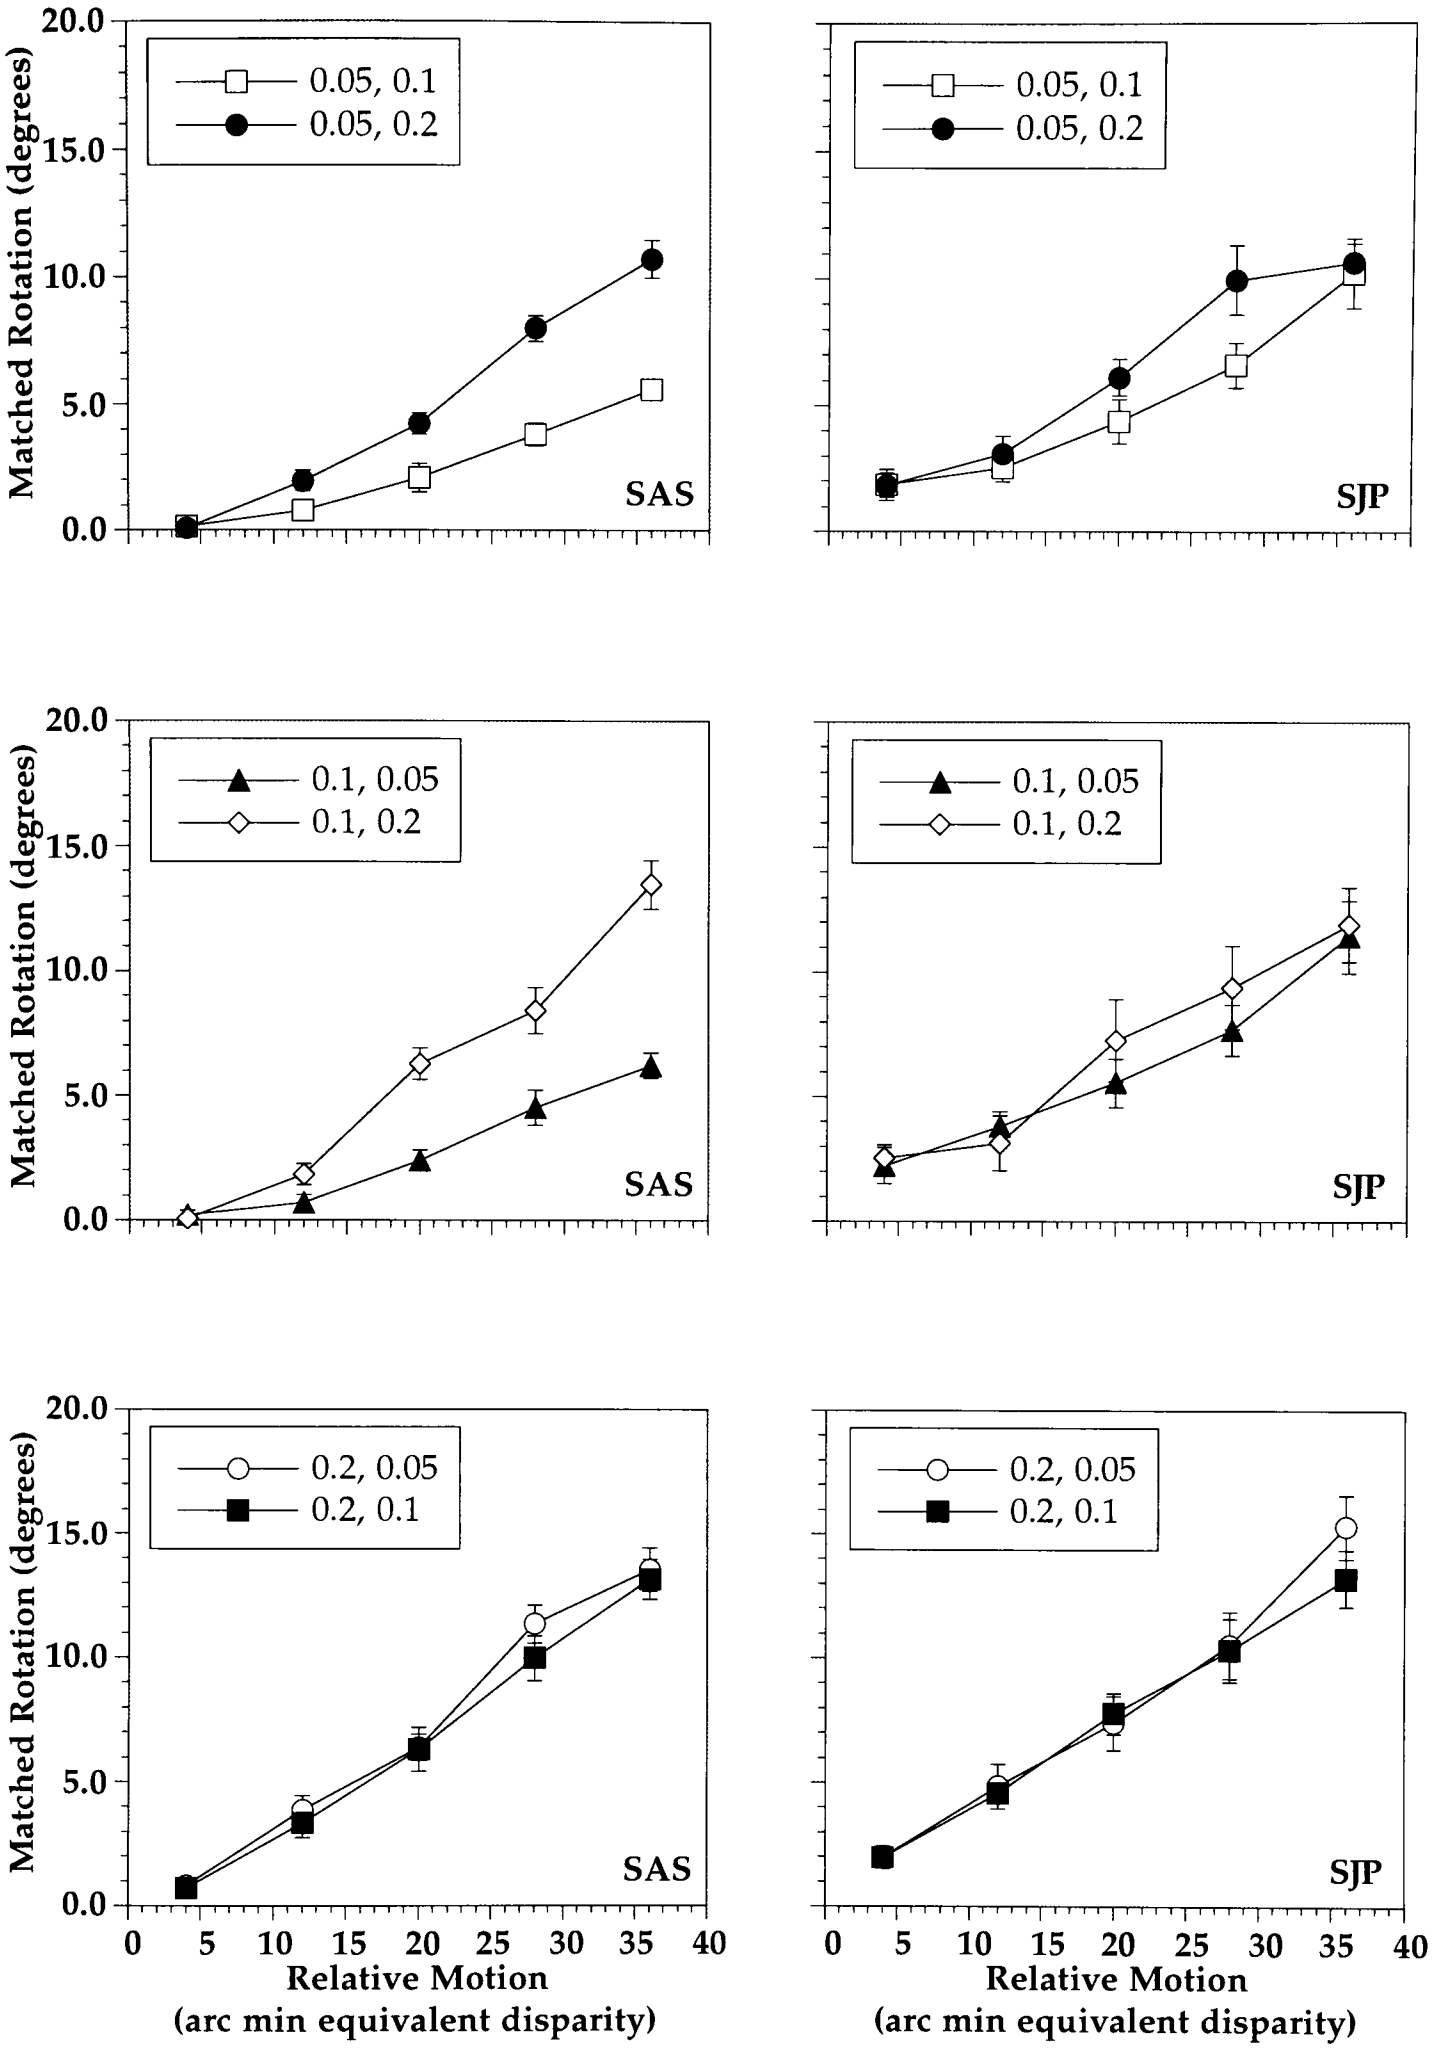
There was a significant main effect of the central corrugation frequency for both subjects ($F[2, 18] = 81.18$, $p<0.0005$; $F[2, 18] = 8.40$, $p<0.01$, for Subjects SAS and SJP respectively). There was a significant main effect of the flanking corrugation frequency for Subject SAS ($F[1, 9] = 26.46$, $p<0.005$) but not for Subject SJP ($F[1, 9] = 1.64$, n.s.). There was a significant main effect of relative motion amplitude for both subjects ($F[4, 36] = 347.94$, $p<0.0005$; $F[3, 36] = 81.17$, $p<0.0005$, for Subjects SAS and SJP respectively).

The two-way interaction between central corrugation frequency and flanking corrugation frequency was significant for Subject SAS ($F[2, 18] = 29.13$, $p<0.0005$) but was non-significant for Subject SJP ($F[2, 18] = 1.96$, n.s.). The two-way interaction between central corrugation frequency and relative motion amplitude was significant for Subject SAS ($F[8, 72] = 5.16$, $p<0.0005$) but was non-significant for Subject SJP ($F[8, 72] = 0.93$, n.s.). The two-way interaction between flanking corrugation frequency and relative motion amplitude was significant for Subject SAS ($F[4, 36] = 12.08$, $p<0.0005$) but was non-significant for Subject SJP ($F[4, 36] = 0.96$, n.s.). The three-way interaction between central corrugation frequency, flanking corrugation frequency and relative motion amplitude was significant for Subject SAS ($F[8, 72] = 5.50$, $p<0.0005$) but was non-significant for Subject SJP ($F[8, 72] = 0.43$, n.s.).

Table 4.1 summarises the significance of each of the main effects and interactions for each subject. A tick indicates that the effect was significant at the 0.05 level or lower. A cross indicates that the effect was non-significant at the 0.05 level.

	Subject	
	SAS	SJP
central c. freq. (0.05, 0.1, 0.2 cpd)	✓	✓
flanking c. freq. (low, high)	✓	✖
relative motion amplitude (4, 12, 20, 28, 36 ameqd)	✓	✓
central c. freq. x flanking c. freq.	✓	✖
central c. freq. x relative motion amplitude	✓	✖
flanking c. freq. x relative motion amplitude	✓	✖
central c. freq. x flanking c. freq. x relative motion amplitude	✓	✖

Table 4.1: Summary of the significance of the main effects and interactions of the above analyses (see Figs. 4.46 - 4.51). A tick indicates that the effect was significant at the 0.05 level or lower. A cross indicates that the effect was non-significant at the 0.05 level.



Figs. 4.46 - 4.51: Rotation Match as a function of Relative Motion for stimuli containing identical central frequencies but different flanking frequencies, for Subjects SAS and SJP.

Figs. 4.52 - 4.57 replot the data to show Rotation Matches for stimuli containing identical flanking corrugation frequencies but different central frequencies. Stimuli with the higher central frequency result in greater rotation matches. Once again, the effect diminished as the common corrugation frequency (i.e. the corrugation frequency present in both stimuli) increased.

Data from each subject were entered into a three-way repeated-measures ANOVA, with flanking corrugation frequency (0.05, 0.1, 0.2 cpd), central corrugation frequency (high, low) and relative motion amplitude (4, 12, 20, 28 and 36 arc min equivalent disparity) as the factors.

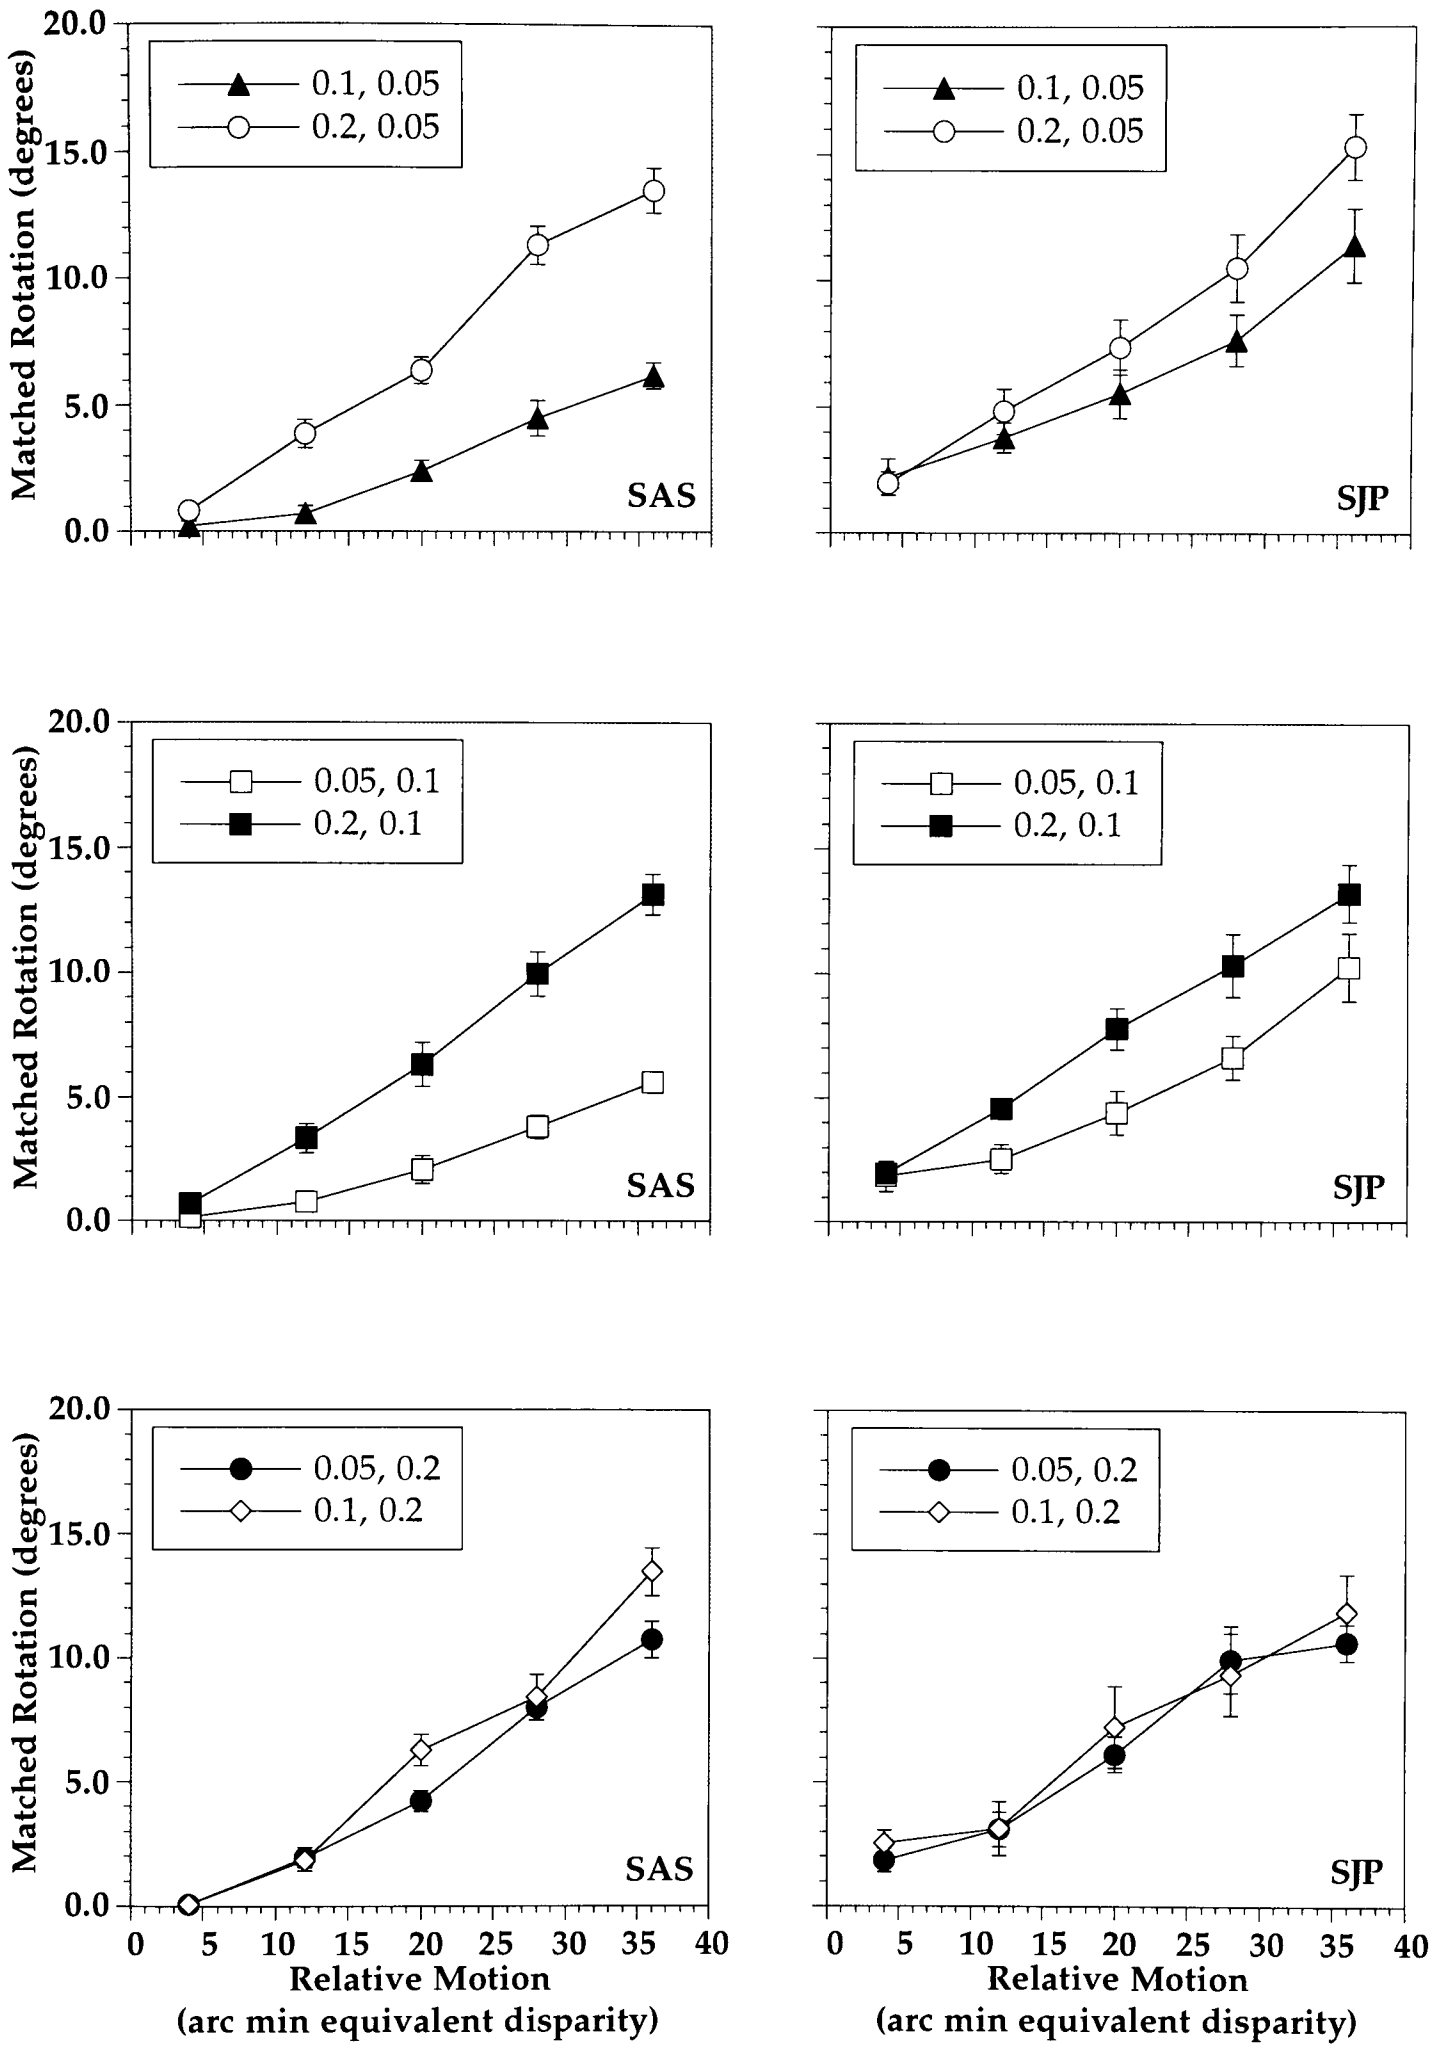
There was a significant main effect of the flanking corrugation frequency for Subject SAS ($F[2, 18] = 5.29$, $p<0.05$) but not for Subject SJP ($F[2, 18] = 1.10$, n.s.). There was a significant main effect of central corrugation frequency for both subjects ($F[1, 9] = 145.77$, $p<0.0005$; $F[1, 9] = 16.17$, $p<0.005$, for Subjects SAS and SJP respectively). There was a significant main effect of relative motion amplitude for both subjects ($F[4, 36] = 347.94$, $p<0.0005$; $F[4, 36] = 81.17$, $p<0.0005$, for Subjects SAS and SJP respectively).

The two-way interaction between flanking corrugation frequency and central corrugation frequency was significant for Subject SAS ($F[2, 18] = 19.79$, $p<0.0005$) but was non-significant for Subject SJP ($F[2, 18] = 2.72$, n.s.). The two-way interaction between flanking corrugation frequency and relative motion amplitude was significant for Subject SAS ($F[8, 72] = 3.10$, $p<0.01$) but was non-significant for Subject SJP ($F[8, 72] = 0.82$, n.s.). The two-way interaction between central corrugation frequency and relative motion amplitude was significant for Subject SAS ($F[4, 36] = 19.05$, $p<0.0005$) but was non-significant for Subject SJP ($F[4, 36] = 0.89$, n.s.). The three-way interaction between central corrugation frequency, flanking corrugation frequency and relative motion amplitude was significant for Subject SAS ($F[8, 72] = 2.76$, $p<0.05$) but was non-significant for Subject SJP ($F[8, 72] = 0.64$, n.s.).

Table 4.2 summarises the significance of each of the main effects and interactions for each subject.

	Subject	
	SAS	SJP
flanking c. freq. (0.05, 0.1, 0.2 cpd)	✓	✖
central c. freq. (low, high)	✓	✓
relative motion amplitude (4, 12, 20, 28, 36 ameqd)	✓	✓
flanking c. freq. x central c. freq.	✓	✖
flanking c. freq. x relative motion amplitude	✓	✖
central c. freq. x relative motion amplitude	✓	✖
flanking c. freq. x central c. freq. x relative motion amplitude	✓	✖

Table 4.2: Summary of the significance of the main effects and interactions of the above analyses (see Figs. 4.52 - 4.57). A tick indicates that the effect was significant at the 0.05 level or lower. A cross indicates that the effect was non-significant at the 0.05 level.



Figs. 4.52 - 4.57: Rotation Match as a function of Relative Motion for stimuli containing identical flanking frequencies but different central frequencies, for Subjects SAS and SJP.

The data are again replotted in Figs. 4.58 - 4.63, to show the effect of reversing the spatial locations of two different corrugation frequencies. These graphs illustrate that both subjects perceive a slightly greater magnitude of rotation in surfaces in which the higher corrugation frequency is in the centre and the lower frequency is in the flanking regions than vice versa.

The data from each subject were entered into a three-way repeated-measures ANOVA, with central corrugation frequency (low, high), corrugation frequency pair (0.05, 0.1 cpd; 0.05, 0.2 cpd; 0.1, 0.2 cpd) and relative motion amplitude (4, 12, 20, 28 and 36 arc min equivalent disparity) as the factors.

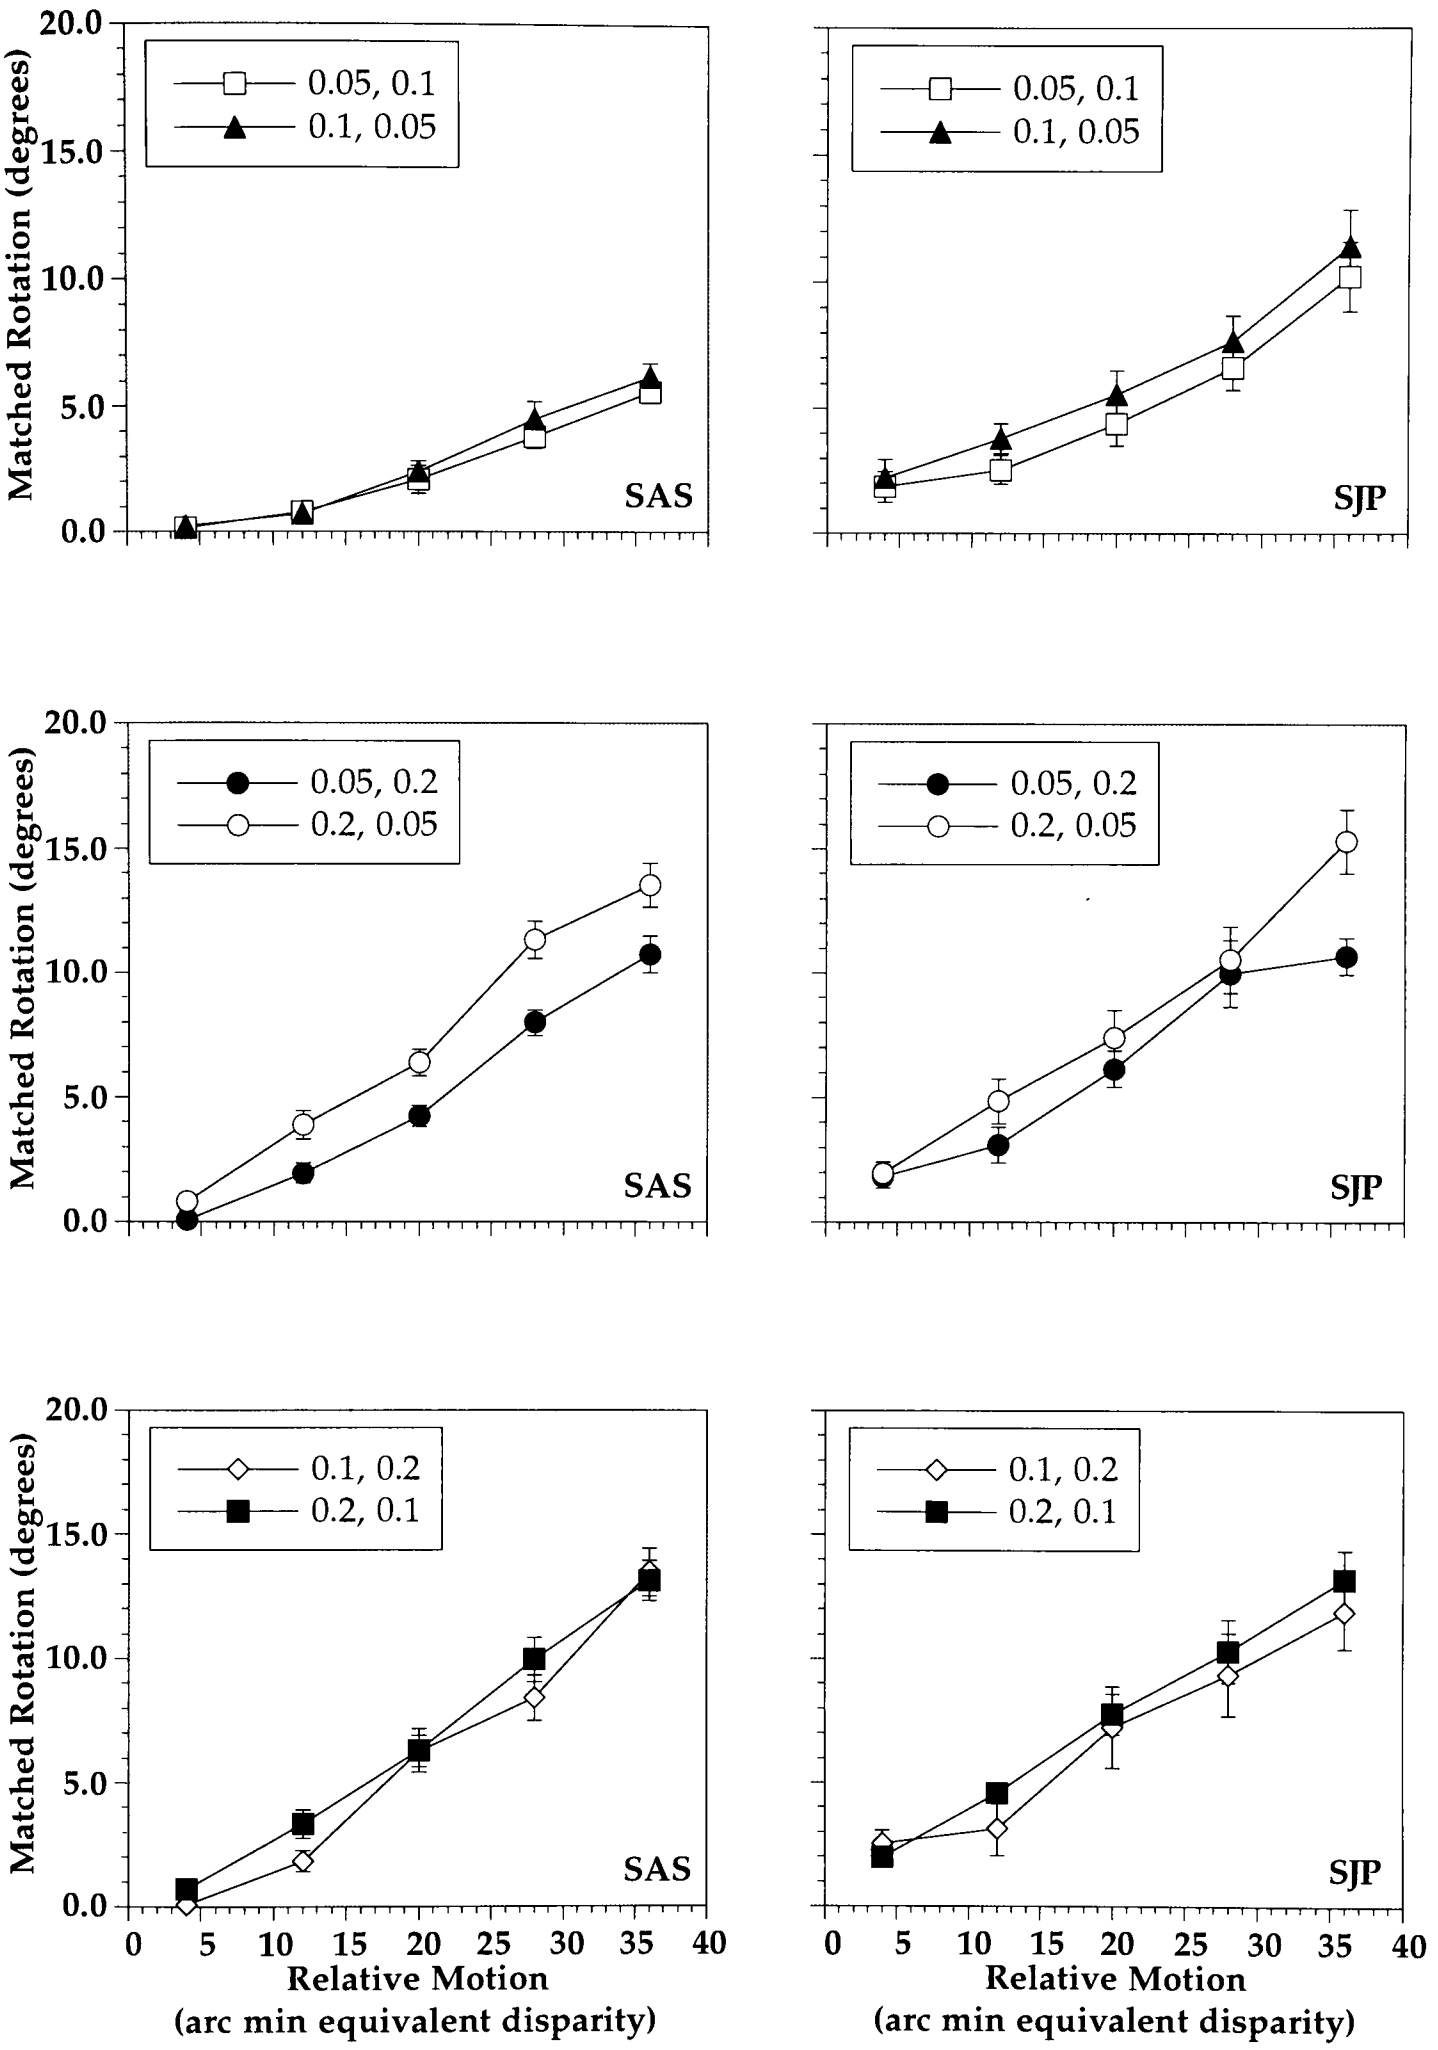
The main effect of central corrugation frequency was significant for both subjects ($F[1, 9] = 31.62$, $p<0.0005$; $F[1, 9] = 9.11$, $p<0.05$, for Subjects SAS and SJP respectively). The main effect of corrugation frequency pair was significant for both subjects ($F[2, 18] = 62.28$, $p<0.0005$; $F[2, 18] = 7.97$, $p<0.01$, for Subjects SAS and SJP respectively). The main effect of relative motion amplitude was significant for both subjects ($F[4, 36] = 347.94$, $p<0.0005$; $F[4, 36] = 81.17$, $p<0.0005$, for Subjects SAS and SJP respectively).

The two-way interaction between central corrugation frequency and corrugation frequency pair was significant for Subject SAS ($F[2, 18] = 8.44$, $p<0.01$) but was non-significant for Subject SJP ($F[2, 18] = 0.48$, n.s.). The two-way interaction between central corrugation frequency and relative motion amplitude was non-significant for both subjects ($F[4, 36] = 0.79$, n.s.; $F[4, 36] = 1.01$, n.s., for Subjects SAS and SJP respectively). The two-way interaction between corrugation frequency pair and relative motion amplitude was significant for Subject SAS ($F[8, 72] = 16.91$, $p<0.0005$) but was non-significant for Subject SJP ($F[8, 72] = 0.94$, n.s.). The three-way interaction between central corrugation frequency, corrugation frequency pair and relative motion amplitude was non-significant for both subjects ($F[8, 72] = 0.98$, n.s.; $F[8, 72] = 0.45$, n.s., for Subjects SAS and SJP respectively).

Table 4.3 summarises the significance of each of the main effects and interactions for each subject.

	Subject	
	SAS	SJP
central c. freq. (low, high)	✓	✓
c. freq. pair (0.05 & 0.1; 0.05 & 0.2; 0.1 & 0.2 cpd)	✓	✓
relative motion amplitude (4, 12, 20, 28, 36 ameqd)	✓	✓
central c. freq. x c. freq. pair	✓	✖
central c. freq. x relative motion amplitude	✖	✖
c. freq. pair x relative motion amplitude	✓	✖
central c. freq. x c. freq. pair x relative motion amplitude	✖	✖

Table 4.3: Summary of the significance of the main effects and interactions of the above analyses (see Figs. 4.58 - 4.63). A tick indicates that the effect was significant at the 0.05 level or lower. A cross indicates that the effect was non-significant at the 0.05 level.



Figs. 4.58 - 4.63: Rotation Match as a function of Relative Motion for stimuli containing identical corrugation frequencies in spatially reversed locations, for Subjects SAS and SJP.

4.7.5 Discussion

The experiment confirmed, in the majority of cases, the prediction that for surfaces containing two corrugation frequencies, perceived rotation is higher for the surface containing the higher mean frequency. The only situation in which this claim was *not* statistically supported was for Subject SJP in the condition in which the flanking corrugation frequency was varied while the central corrugation frequency remained constant.

It was also shown that the size of this effect was dependent upon both the location and the value of the common corrugation frequency. The effect was larger (and significant for both subjects) when the common frequency was in the flanking portions of the surface rather than in the central region, when the effect was significant for Subject SAS only — meaning that variations in the central frequency have a greater effect upon perceived rotation than variations in the flanking frequency. The relative importance of the central region is also demonstrated by the finding that the perceived rotation of surfaces with identical corrugation frequencies is greater when the higher frequency is in the central, rather than the flanking, position.

4.8 Summary and Conclusions

The experiments of this chapter examined the influence of a variety of factors on the magnitude of perceived rotation, extending the results of the previous chapter by examining perceived rotation at relative motion values greater than the Transition Point. Experiment 6 showed that stimuli of different corrugation frequencies appear to rotate to the same extent when the surface of lower frequency contains a greater amplitude of relative motion than the surface of higher frequency. Experiment 7a showed that the perceived magnitude of rotation increases with increasing corrugation frequency, a finding in qualitative accordance with the results of Experiments 1a and 6. The data of Experiment 7a indicate, however, that perceived rotation is *not* determined solely by the MMG, even for the frequency range within which the MMG largely determines the Transition Point. This conclusion is supported by the results of Experiment 7b, which used peak-clipped stimuli to examine the influence of the MMG on the

perceived rotation of surfaces of three different corrugation frequencies. This experiment found firstly that the 0.05 cpd surface was the only surface for which the MMG appreciably affected perceived rotation. The perceived magnitudes of rotation of the 0.1 and 0.2 cpd surfaces were unaffected by the MMG, and it was suggested that for surfaces of higher corrugation frequency there is a ceiling level of rotation which may be perceived for a given amount of relative motion.

Secondly, it was found that stimuli having identical MMGs were not perceived to rotate to the same degree: surfaces of higher corrugation frequency (and thus a lower level of 'clippedness') were perceived to rotate more, especially at higher relative motion amplitudes.

Both of these findings differ from those of Experiments 1a and 1b, which showed firstly that the Transition Point decreased with increases in the MMG for surfaces of all three corrugation frequencies of Experiment 7b, and secondly that the Transition Points of stimuli containing identical MMGs were approximately equal.

It was suggested that the reason for the ceiling level of rotation in the stimuli of higher corrugation frequency was a trade-off between rotation and rigidity (Williams, 1993). The reason for the finding that surfaces of high corrugation frequency were perceived to rotate to a greater extent than those of identical MMGs but lower corrugation frequency, may be the similarity of the low frequency, peak-clipped stimuli to square waves, which are generally perceived to shear than to rotate (Ono and Steinbach, 1990).

Experiment 8 found that perceived rotation is affected by the phase of the surface, but that this dependence decreases with increasing corrugation frequency. Surfaces of 0.05 cpd which contained a central peak were perceived to rotate to a considerably greater degree than those containing a central trough; this trend was reversed for 0.1 cpd surfaces and was absent for 0.2 cpd surfaces. It was suggested that the determinant of the relative magnitudes of perceived rotation is not the form of the central corrugation *per se* but is rather the presence of adjacent corrugations. The qualitative difference between results obtained with the 0.05 cpd surface and those obtained with the 0.1 and 0.2 cpd surfaces remains unclear.

Experiment 9 showed that perceived rotation decreases with increasing stimulus size. Two factors were suggested to account for this trend: firstly, the increase in the vertical perspective information signalling zero rotation, and secondly the increase in the extent to which the surface may be perceived to be 'fixed' in the plane of the monitor and thus constrained by the monitor's own lack of rotation.

The increase in perceived rotation with decreasing stimulus size was found to be greater than the corresponding decrease in Transition Point (Experiment 3). For each halving of stimulus size the increase in perceived rotation was approximately 0.35, whereas Transition Points increased by a factor of 0.2 for each doubling of stimulus size. This indicates that vertical perspective plays a larger role in the perception of rotation when the percept is at supra-threshold, rather than threshold, levels. This result contrasts with the relative strengths of the MMG in the determination of supra-threshold rotation and the Transition Point. The perception of supra-threshold rotation is much less heavily influenced by the MMG than is the Transition Point.

Using surfaces comprising two different corrugation frequencies, Experiment 10 found that the relative perceived rotation of two surfaces containing a common corrugation frequency of either 0.05 or 0.1 cpd was generally dependent on both the value and the spatial location of the second frequency. Surfaces containing a common corrugation frequency of 0.2 cpd, however, were perceived to rotate to a more similar extent. The experiment also showed that the magnitude of perceived rotation was dependent upon the spatial locations of the two corrugation frequencies: pairs of surfaces containing identical frequencies in reversed spatial locations were *not* perceived to rotate to the same extent. Subjects perceived a greater amount of rotation in the surface in which the higher corrugation frequency was in the central position than in the surface in which it was in the flanking position, indicating the greater importance of the central region in the determination of perceived rotation magnitude.

Chapter 5:

The Trade-Off between the Perceived Magnitudes of Depth and Rotation in Observer-Produced Motion Parallax Surfaces

5.1 Introduction

In Section 1.6.2 it was noted that the perceived quantities of two qualitatively different attributes of a motion parallax stimulus may exist in a trade-off relationship. Trade-offs have been noted in a variety of stimuli types. An increase in the perceived magnitude of one attribute goes hand in hand with a decrease in the perceived magnitude of the other. One of the attributes is invariably related to the structure of the object — its depth or its shape for example. The other attribute is related to the object's motion — for example the extent of rotation, or the magnitude of angular velocity. Numerous studies have examined the concomitant variations in such attributes.

5.1.1 Work involving non-quantitative judgements

Early work attempting to determine the relationship between perceived rotation and depth involved non-quantitative judgements. Braunstein (1966) presented subjects with spheres which rotated about either the X or the Y axis. Projection was either polar or parallel. Observers selected from a set of real, 3-D models the one which most resembled the observed display. They mounted the model on a rotation device, and used a toggle switch to make the object translate/rotate in a given direction. It was found that judgements of 3-D structure were nearly identical for the two projection types, but that rotation direction '... was correctly judged only in the case of the polar projections'. Braunstein (1976) argued that this was evidence for the separability of processes involved in depth and

rotation judgements. This claim is not necessarily warranted, however. Braunstein's (1966) 3-D judgements were not concerned with the depth order of the stimulus, but merely with its form (subjects decided whether the simulated structure was a circle, sphere or a cylinder). The fact that parallel and perspective projections gave rise to the same perceived 3-D form is relatively unsurprising and somewhat irrelevant to the question being addressed. The fact that rotation judgements were different in the two projection conditions is consistent with the fact that rotation direction and depth order are undetermined in the parallel case but not in the perspective case. There is no 'incorrect' perceived direction of rotation or depth order for the former projection type. The percepts are, however, theoretically interdependent. That is, a reversal in one percept occurs concomitantly with a reversal in the other. Therefore in order for Braunstein's (1976) claim to hold, it would have to be demonstrated that for a given projection type the perceived depth order was independent of perceived rotation direction.

In a later attempt to examine the possible dependence of depth and rotation judgements, Petersik (1980) showed observers spheres which rotated stroboscopically about the vertical axis. Frame duration, interstimulus interval (ISI), sphere size, simulated projection distance, dot numerosity and number of frames per 360° rotation were manipulated. Subjects rated the depth of the sphere and its apparent direction of rotation. It was found that these were related to a given independent variable in different ways, suggesting a degree of independence of depth and rotation. It was also found, however, that the magnitude of perceived depth was dependent on the correctness of perceived rotation direction, implying that the two judgements were not completely dissociated. In a later paper, Petersik (1991) investigated whether information regarding rotation magnitude develops prior to, in parallel with, or after information regarding the recovery of structure. His results were inconclusive, but he argued in favour of the third alternative.

Braunstein and Andersen (1981) presented observers with a translating, horizontally-oriented dihedral angle (see Fig. 1.5). It was found that the structures were often perceived to rotate about a vertical axis as they translated. These rotating structures were divisible into two groups. The first group contained structures which were reversed in depth, with the

slowly moving dots perceived as nearer than the quickly moving dots. In this case, the display roughly approximated a configuration in which the faster dots rotated about the slower dots. The authors observed that a small amount of non-rigidity must be either tolerated by the visual system (a suggestion later confirmed by Hogervorst et al., 1996) or perceived by the observer. The second group of rotating structures contained surfaces with perceived depth orders identical to those simulated. The authors suggested that the perception of rotation of such structures might occur with a concomitant underestimation of surface slant (i.e. an underestimation of depth). Such measurements were not made, however. Braunstein and Tittle (1988) also suggested that apparent rotation may arise when subjects perceive less depth than that simulated by the velocity gradient and some of the gradient is therefore attributed to the effects of rotation. Again, however, simultaneous measurements of perceived depth and rotation were not made.

5.1.2 Work involving quantitative judgements

Braunstein, Liter and Tittle (1993) presented subjects with the dihedral angle described above. The structure either translated horizontally under perspective projection (Fig. 5.1: a, b) or rotated under parallel projection (Fig. 5.1: c, d). Perceived depth order (centre-near or centre-far) and angle magnitude judgements (β) were made by adjusting the direction and magnitude of a comparison angle.

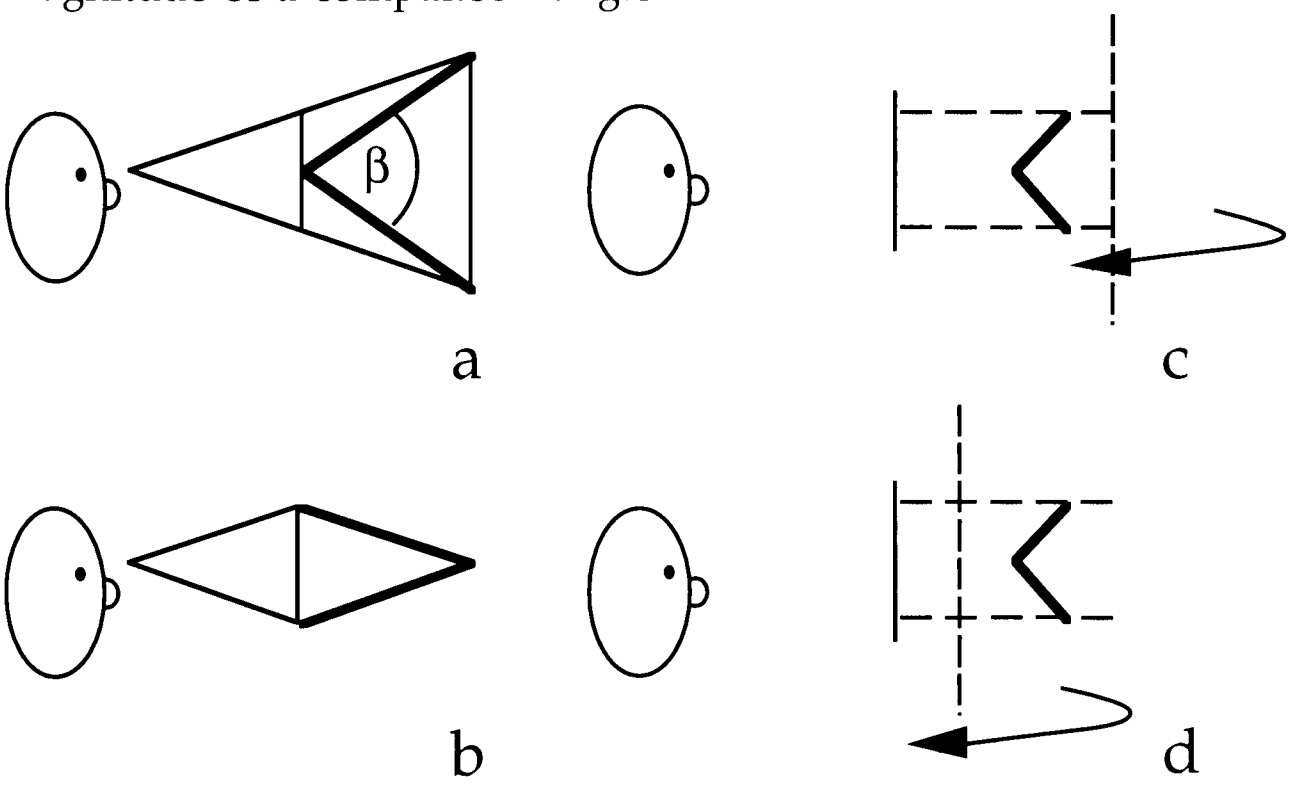


Fig 5.1: Representation of the dihedral angle stimuli used by Braunstein and Tittle (1993). Taken from Braunstein and Tittle (1993, Fig. 1, p. 599).

Observers tended to over-estimate the magnitude of the dihedral angle when the structure was under perspective projection, and to under-estimate the magnitude when it was under parallel projection. Additionally, in accordance with Braunstein and Andersen (1981), the perceived depth order of the translating structures was sometimes reversed. On the basis of these findings it was suggested that observers were analysing the structures using a compromise between perspective and parallel analyses. However, an alternative possibility, namely that the displays were perceived as perspective projections of objects that were both rotating and translating, was also considered. It was hypothesised that the difference in response to the translations and the rotations was based upon the horizontal compression of the flow field which occurs during rotation. Braunstein et al.'s Experiment 5 (p. 608) introduced compression into the displays, and found that the judged dihedral angle increased (i.e. perceived depth decreased) with increasing amounts of compression. Despite the fact that rotation judgements were not taken it seems likely that subjects perceived a greater amount of rotation with the increased compression, and that a trade-off between depth and rotation occurred.

This was confirmed by Liter and Braunstein (1994), who compared rotation and depth judgements for sequences in which the frontal view was either present or absent. Perceived rotation was greater when the frontal view was not in the sequence. The dihedral angle was judged as flatter in these conditions (i.e. perceived depth was less). These results are consistent with a depth-rotation trade-off.

Van Veen and Werkhoven (1996) investigated the relationship between the perceived rotation and the perceived slant of a plane. The display consisted of a test plane and a match plane, under parallel projection. Both stimuli were slanted about their central horizontal axis, and both rotated about their central vertical axis. The slant magnitudes were 15, 30, 45 and 60°. The rotation magnitudes were 28, 42, 70 and 98°. The task of the observer was to match the amount of perceived slant and perceived rotation of the test stimulus by adjusting the slant and the rotation of the match stimulus.

It was found that over a range of conditions, perceived slant and perceived rotation were highly negatively correlated, indicating that a

trade-off between the two attributes occurred. Within other conditions these attributes were less highly correlated. Generally, it was the case that at low rotation magnitudes, subjects' slant/rotation judgements were highly negatively correlated throughout the range of slant magnitudes. As the magnitude of rotation increased the slant/rotation judgements became decreasingly correlated, such that at the highest rotation magnitude (98°) and lowest slant magnitude (15°) the correlation had almost disappeared. The main conclusion that can be drawn from these data, however, is that the trade-off between depth and rotation occurs over a wide range of simulated conditions.

Evidence for an association between perceived depth and rotation has also been found in studies on the recovery of shape from deforming contours. Cortese and Andersen (1991) displayed simulated silhouettes of ellipsoids rotating about a vertical axis. The ambiguity of the projected image is illustrated below (Fig. 5.2).

In Experiment 3 (p. 323) the major axis of the object and the extent of simulated rotation were manipulated. Subjects made judgements of both of these, and it was found that the shape responses could be accurately predicted from the rotation responses. An under-estimation of rotation resulted in an over-estimation of the depth. Likewise, an over-estimation of rotation resulted in an under-estimation of the depth. The authors suggested that the extent of rotation is recovered through a comparison of the maximum and minimum horizontal extents of the image contour, a ratio which accounted for 94% of the variance in the rotation responses, and that from this estimate the shape could be recovered.

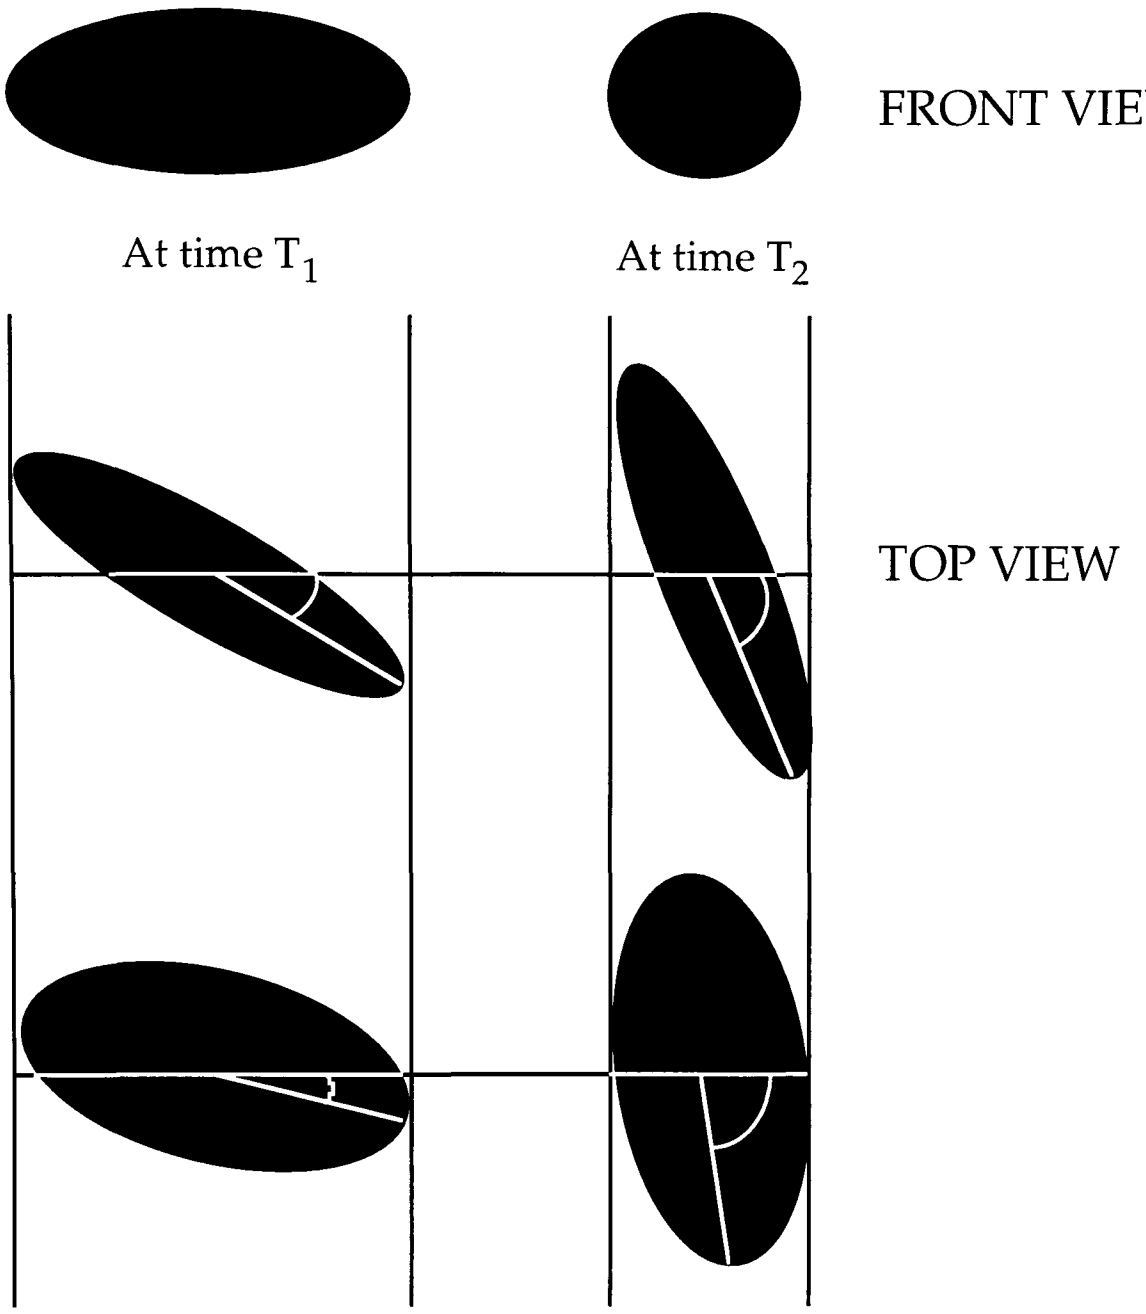


Fig 5.2: Representation of the rotating ellipsoids of Cortese and Andersen (1993). Taken from Cortese and Andersen (1993, Fig. 2, p. 316).

Pollick (1987) showed that rotating ellipsoids which were symmetric about the axis of rotation were judged to contain less depth than ellipsoids which were not symmetric. Pollick and Wilson (1991) presented both types of ellipsoid and asked observers to judge their angular velocity and shape. It was found that ellipsoids symmetric about the axis of rotation were judged as containing less depth and having a greater angular velocity than non-symmetric ellipsoids. Pollick (1994) also showed that ellipses which rotated about a non-vertical axis were judged to rotate more slowly and to contain a greater amount of depth than those which rotated about a vertical axis.

A further illustration of the trade-off between shape and angular velocity is provided by Pollick (1997). Observers were presented with orthographically-projected structures which underwent uniform affine

stretches along the line of sight. A discussion of affine transformations is not relevant here, suffice to say that such transformations do not provide sufficient information for the recovery of absolute distances and angles between parts of an object — Euclidean metric structure. They provide information only up to a stretching transformation along the line of sight. Norman and Todd (1993) had earlier used stimuli of this nature and found that affine stretches along the line of sight were invisible to observers. They interpreted this result as supporting their earlier claims that the visual system performs an affine, rather than a Euclidean, analysis of SFM stimuli (e.g. Todd and Bressan, 1990; Todd and Norman, 1991). However, Hogervorst et al. (1996) presented subjects with the same stimuli and found that the structures were perceived as rigid. The authors presented a model which demonstrated that there was a near-rigid Euclidean interpretation of the stimulus, in which the 2-D motion arising from the affine stretch could alternatively be interpreted as a change in the angular velocity. Pollick (1997) asked observers to monitor changes in perceived angular velocity, and found that the occurrence of such changes was highly correlated with the occurrence of the simulated affine stretching, thus illustrating a trade-off between depth and angular velocity.

One set of experiments did not find evidence for a trade-off between depth and rotation. Liter, Braunstein and Hoffman (1992, 1993) investigated the perceived depth magnitude in 2- and 30-view apparent motion displays. Stimuli were orthographic projections of a structure undergoing rotation about a fixed axis. Each structure contained five dots. In the first three experiments, the task of the subject was to pick from a 'response set' the 3-D structure which most closely resembled that shown. In the fourth experiment the subject adjusted the distance between two dots to match the perceived depth of the structure along the line of sight. The results indicated that judged depth increased when either the simulated depth *or* the simulated rotation increased. Furthermore, in an unpublished study in which depth and rotation judgements were made on the same trial, no evidence for a trade-off relation was found. This led to the conclusion that perceived depth was determined by the amount of relative motion, regardless of its source. The authors proposed that perceived depth may be calculated as follows. The visual system extracts components of image velocities which are not due to curl. Removal of this component has been suggested in theories of recovery of heading

from optic flow (e.g. Koenderink and van Doorn, 1981; Longuet-Higgins and Prazdny, 1980; Perrone, 1992). The pair of points with the maximum signed difference in velocity is selected from the remaining points, and a measure of perceived depth is calculated from the relative velocities of these.

These results are consistent with Proffitt, Rock, Hecht and Schubert's (1992) suggestion that relative motion is used as a heuristic for determining perceived depth. This approach has also been proposed in comparisons of depth from disparity and SFM (Durgin, Proffitt, Olson and Reinke, 1995). However, since perceived depth may be affected by other image variables such as display width (Caudek and Proffitt, 1993; Proffit et al., 1992) and compression (Braunstein et al., 1993) — factors which would be expected to affect perceived rotation — it is evident that depth and rotation are not independent. Liter et al. (1993, p. 1463) note that the apparent discrepancy between their results and those of Braunstein et al. (1993) might be due to stimulus differences. The latter study used '... densely textured, planar surfaces which provide richer information for judging rotation than the five-dot displays in the present study'. Also, the investigation with dihedral angles used rotation about a vertical axis whereas the axes used by Liter et al. (1993) were slanted. The authors proposed that processes which recover depth from rotation may use the amount of rotation more effectively when the axis is in the image plane. This suggestion accords with Loomis and Eby (1988), who found that the perceived depth of a SFM sphere was greatest when the axis of rotation was in the image plane and decreased as the axis approached the line of sight. It also fits with the results of Pollick (1994) outlined above.

In a study investigating the ability of the visual system to calibrate motion parallax with absolute distance information, Ono et al. (1986) presented observers with head-linked motion parallax stimuli at distances of 40, 80, 160 or 320 cm. There were three extents of distal parallax, corresponding to equivalent disparities of 28.2, 56.4 or 112.8 arc min. Subjects indicated whether they perceived (i) depth with no motion — an apparent rigid three-dimensional surface without movement; (ii) depth with motion — an apparent rotation of a corrugated surface about a vertical axis; or (iii) no depth at all. If either (i) or (ii) were perceived, observers indicated the magnitude of perceived depth of the surface.

Increases in either viewing distance or parallax motion led to increases in the number of observers perceiving rocking motion. For any given condition, the mean amount of perceived depth was *lower* when the surface was perceived to be rocking than when it was perceived to be stationary.

Ono and Steinbach (1990) presented observers with random dots moving with a velocity gradient of either a sine, triangle, ramp or square waveform. The motion was linked to head motion in one condition and to no head motion in the other. The amplitude of relative motion was either 12 or 113 arc min equivalent disparity. Observers indicated the amount of perceived depth and decided whether the dots appeared to translate. If the dots appeared to translate, subjects chose dots at the 'furthest' and 'nearest' parts of the stimulus and indicated the extent of perceived dot motion on a ruler. It was found that the sine waveform contained the greatest amount of depth and the least amount of motion. The square waveform contained the least amount of depth and the greatest amount of motion. The relative extents of depth and motion in the triangle and the ramp waveforms fell between these. Additionally, for each waveform less depth and more motion were perceived in the no head movement condition than in the head movement condition. The authors interpreted these data as suggestive of a trade-off between depth and motion.

There have been two basic approaches to the processes assumed to underlie the perception of rotation in head-linked parallax stimuli. One approach is basically a re-description of the phenomenon, couched in terms of location constancy (Graham, 1951; Ittelson, 1960; Wallach, 1985; Mack, 1986; Swanston and Wade, 1987). As will be recalled from Section 3.2.6, this is the compensation process which allows a moving observer to perceive the surrounding environment as stationary, despite the retinal motions caused by the observer's own movements. Ono et al. (1986) and Ono and Steinbach (1990) have invoked Wallach's 'immobility range' to explain the perceived rotation of parallax structures.

As outlined above, the 'immobility range' (Wallach, 1985, 1987) refers to the range of externally-produced object motions which are perceived by a moving observer to have arisen from self-motion and which thus result

in the perception of a stationary object. Of the cases examined by Wallach, the most applicable to parallax paradigms is the situation in which an object is made to rotate as the observer moves past it. Considering first a stationary object, observer translation results in an observer-relative rotation of the object which is entirely attributable to the changing position of the observer. If the object is actually made to rotate, the observer-relative rotation now comprises two components. These are firstly the rotation described above — namely the rotation due purely to self-movement — and secondly the rotation due to the independent rotation of the object itself. Wallach (1985) showed that an object could rotate through an angle 40% of its relative rotation before subjects perceived it as objectively rotating.

Ono et al. (1986) suggested that "... rocking motion represents a percept outside the 'immobility range' ...". It is unclear that an unmodified version of Wallach's framework can legitimately be applied to the perceived rotation of a parallax stimulus, however. As used by Wallach in the context specified above, the term 'immobility range' refers to all situations in which the objective rotation of an object is not perceived. There are no such situations in the case of parallax surfaces, since parallax motion is always ambiguous and thus can never be said to arise from an objective rotation.

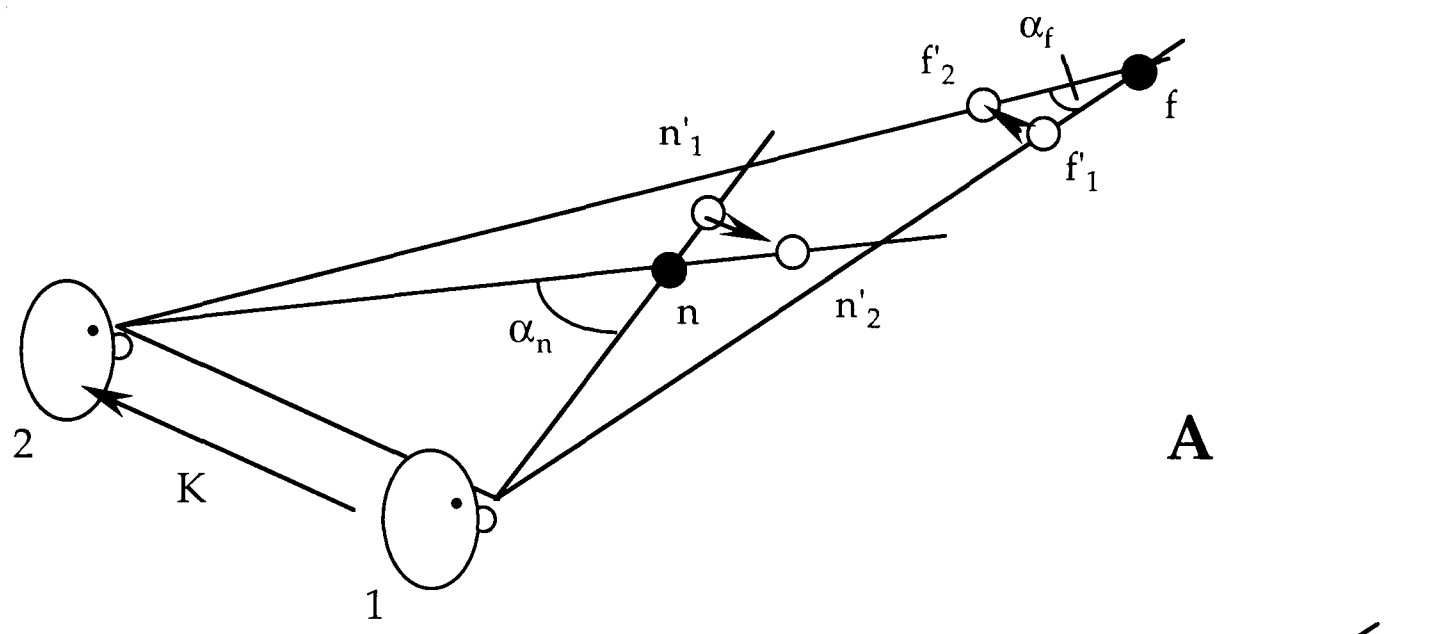
The application of Wallach's framework to this situation would therefore involve an extrapolation from the finding that moving observers are biased towards perceiving an objectively rotating object as stationary. This extrapolation would involve the assumption that if the target is *ambiguous* with respect to its objective rotation (as is the case with parallax surfaces), then the relative biases of a moving and a stationary observer towards a stationary interpretation of the scene might result in a difference in their willingness to interpret a given pattern of motion as arising from a rotation. We might expect rotation to be perceived less willingly when the observer is moving than when they are stationary. This suggestion is supported by the finding that more rotation and less depth was perceived in parallax surfaces viewed by stationary, as opposed to moving, observers (Ono and Steinbach, 1990).

However, the fact remains that the 'immobility range' framework is

merely a loose re-description of the fact that rotation is sometimes perceived in head-linked parallax stimuli. It provides no explanation of *why* rotation should be perceived under some circumstances, and moreover is purely qualitative.

Ono et al. (1986) touched upon a more quantitative approach to their trade-off data in their drawing of parallels between their data and the findings of Gogel (e.g. 1980, 1981). It has been shown that objects sometimes appear to move concomitantly with motion of the head, even when the objects do not physically move (Hay and Sawyer, 1969; Wallach, Yablick and Smith, 1972; Gogel, 1976, 1977, 1979; Gogel and Tietz, 1977). It has been suggested that this apparent motion is determined by retinal motion in conjunction with perceived distance (Gogel and Tietz, 1974, 1977; Gogel, 1976, 1977, 1979). Gogel (1980) provided the diagram below (Fig. 5.3) to illustrate this relationship.

In this figure the observer, while moving their head laterally through a distance K, views two physically stationary points of light, represented by the filled circles, in an otherwise dark field. The nearer light is marked n and the further light is marked f. The degree of relative retinal motion which occurs during the observer's translation is determined by the physical distances of n and f from the eye. Since these distances are identical in the lower and the upper diagrams, the relative retinal motion will also be identical.



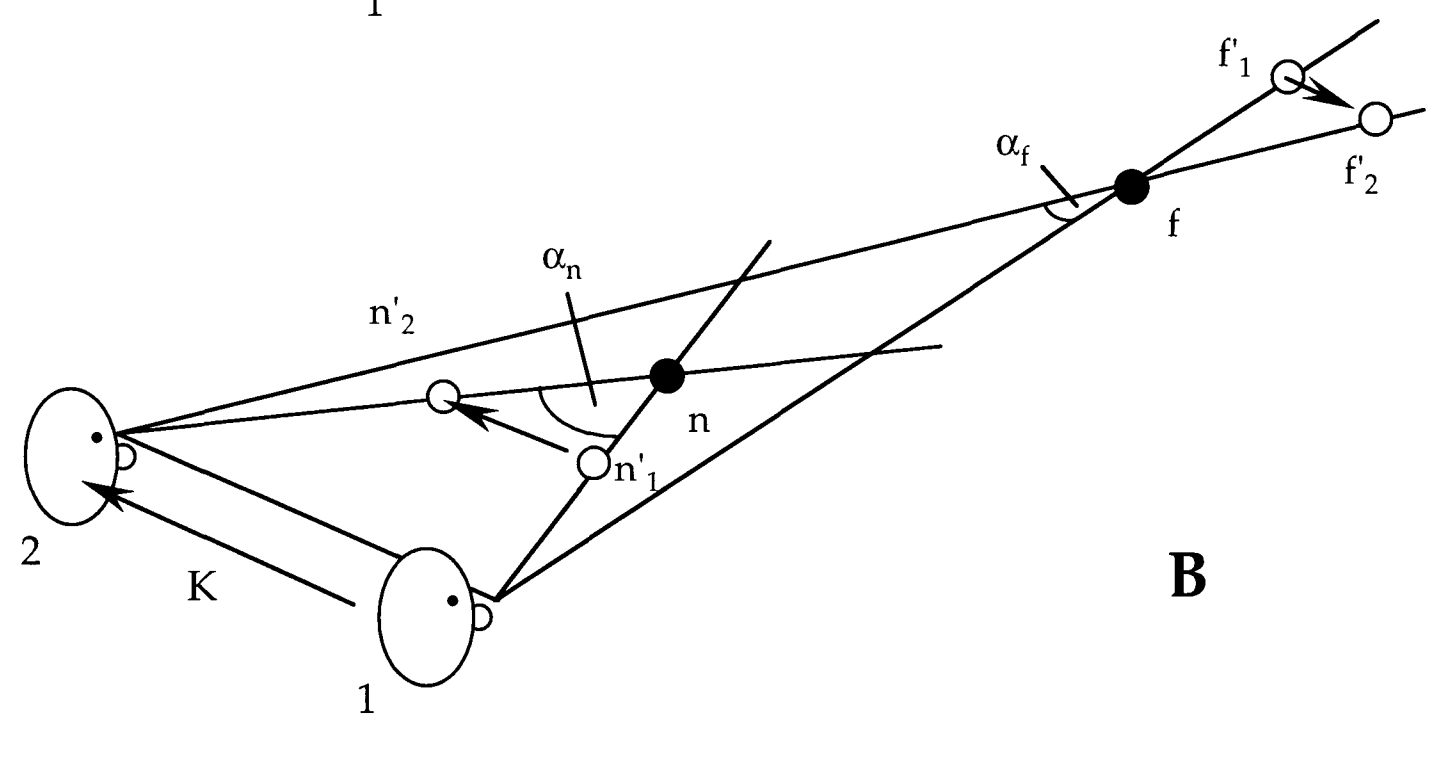


Fig 5.3: Illustration of the relationship between the perceived distance between two points, their perceived motion and the perceived absolute distance. Taken from Gogel (1980, Fig. 1, p. 156).

Three different perceptual situations may occur. In the first, the distances of points n and f are correctly perceived. When the head moves, despite the relative motions of the lights' images, the two points are perceived as stationary. In the second situation (Fig. 5.3A), the depth between n and f is perceptually underestimated: point n is perceived to be further, and point f nearer, than its physical distance. The points will therefore appear to move in opposite directions, from n'_1 to n'_2, and from f'_1 to f'_2 respectively. In the third situation (Fig. 5.3B), the depth between n and f is perceptually overestimated: point n is perceived to be nearer, and point n further, than its physical distance. The points again appear to move in opposite directions. This time, however, each point moves in the opposite direction to the direction in which it moved in the previous situation.

If points n and f are respectively treated as points on a peak and a trough of a parallax surface, then it can be seen that A represents a situation in which counter-rotation (or 'convex' rotation - see Fig. 1.4C) of the surface is perceived at the expense of depth. The greater the extent of perceived counter-rotation, the smaller the amount of perceived depth. B represents a situation in which the magnitude of perceived depth is greater than simulated. The surface rotates in the same direction as the head movement ('concave' rotation - see Fig. 1.4B).

However, despite recognising the similarity of the principles underlying the percepts from Gogel's displays and their own, Ono et al. (1986) did not develop a quantitative approach relating explicitly to their own paradigm.

Rogers and Collett (1989) provided a theoretical basis for a trade-off between depth and rotation in motion parallax stimuli (see Fig. 1.4). The putative trade-off has been observed empirically (e.g. Ono et al., 1986; Ono and Steinbach, 1990). To date, however, no published work using parallax stimuli has presented a model which explicitly relates the expected magnitudes of depth and motion. Such a model is necessary if quantitative, rather than merely qualitative, analyses of stimulus motion are to be made.

In the unpublished thesis of Williams (1993), a geometric model showing permissible combinations of depth and rotation for a given amount of parallax motion was presented. A modified version of that model is presented below (Figs. 5.4 and 5.5). The differences between the present model and that of Williams (1993) are firstly, in the original model the rays of light from points P and Q' were treated as parallel in order to simplify the geometry. Secondly, the motion parallax in the current model is created through observer movement whereas in the original model it was created by stimulus movement.

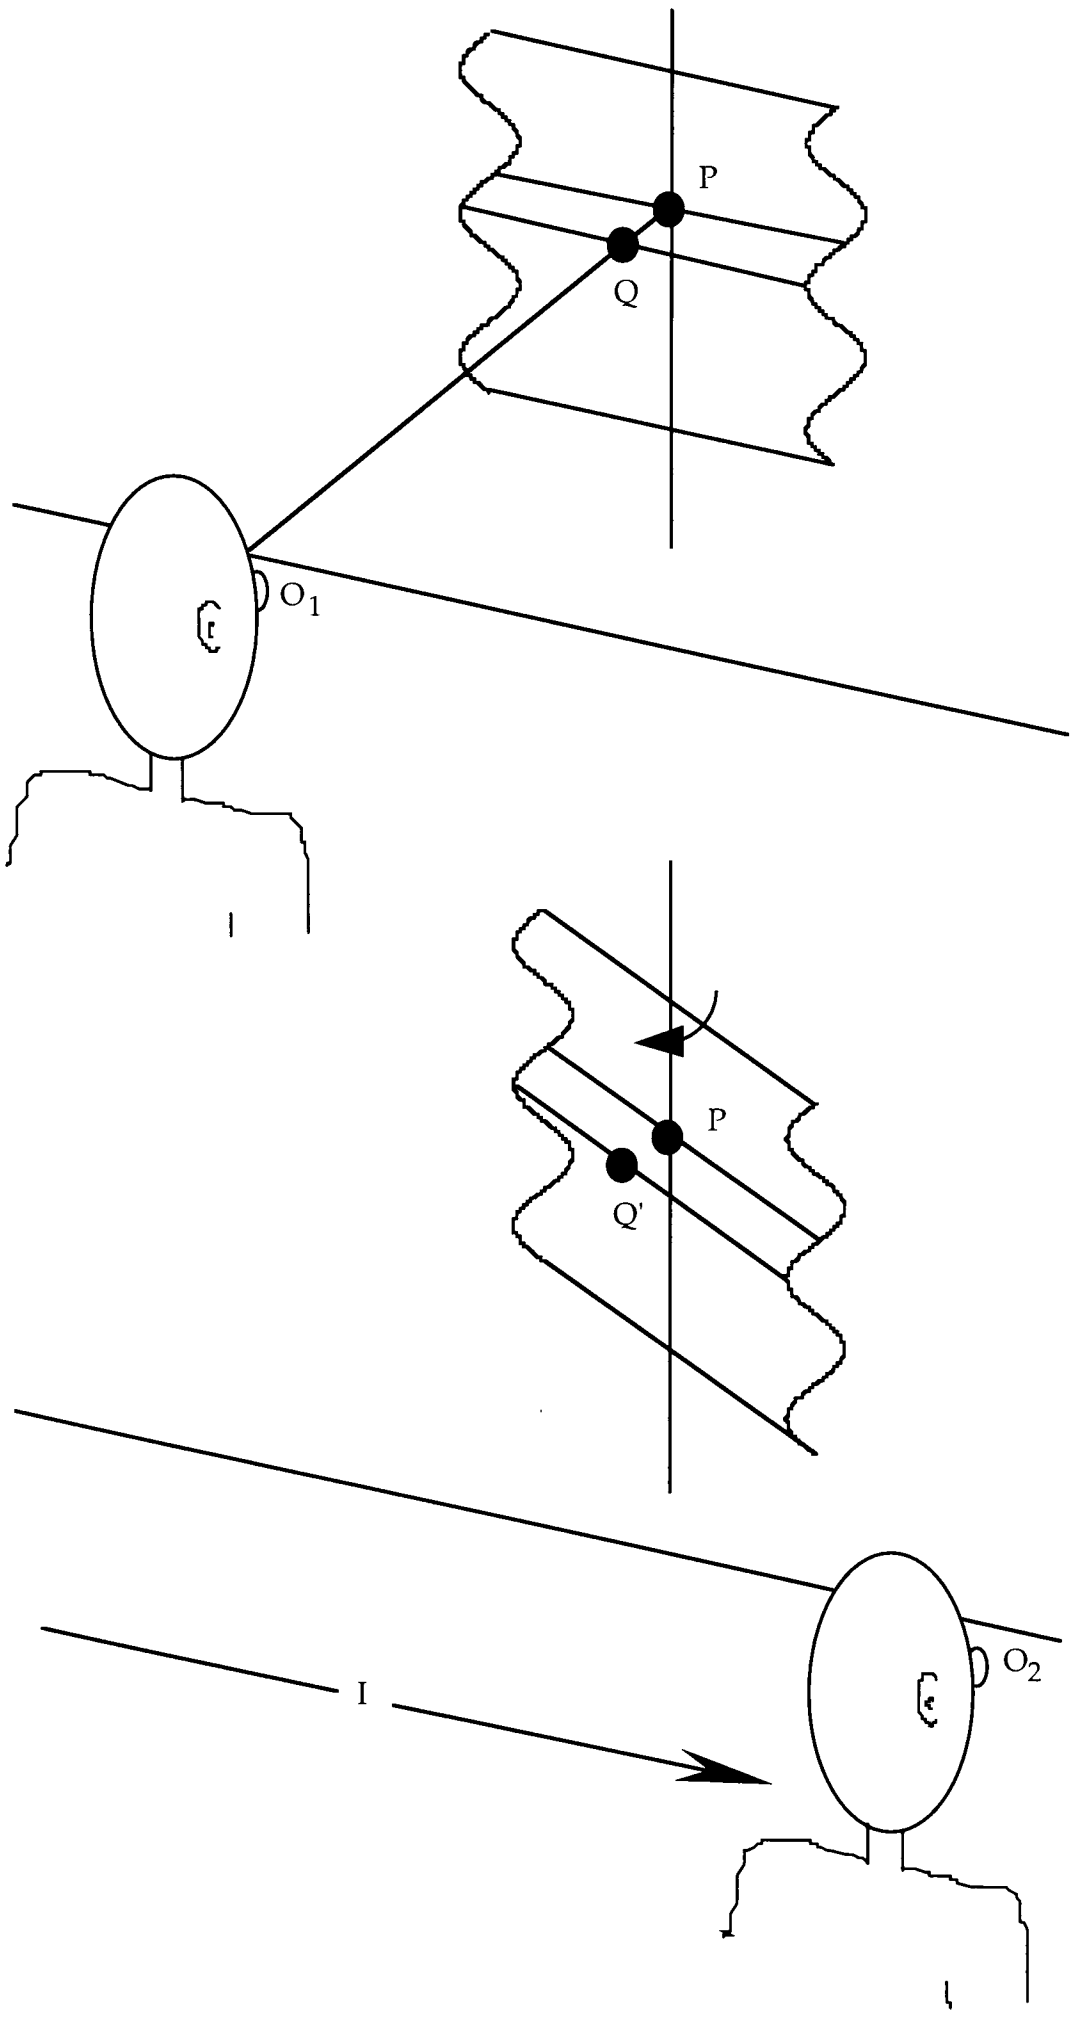


Fig 5.4: A horizontally corrugated sinusoidal surface rotates about a vertical axis through Point P as the eye of the observer moves from O_1 to O_2 through distance I. Point Q (diagram 1) is situated on a peak. Point Q′ (diagram 2) is the new position of Point Q following the structure's rotation.

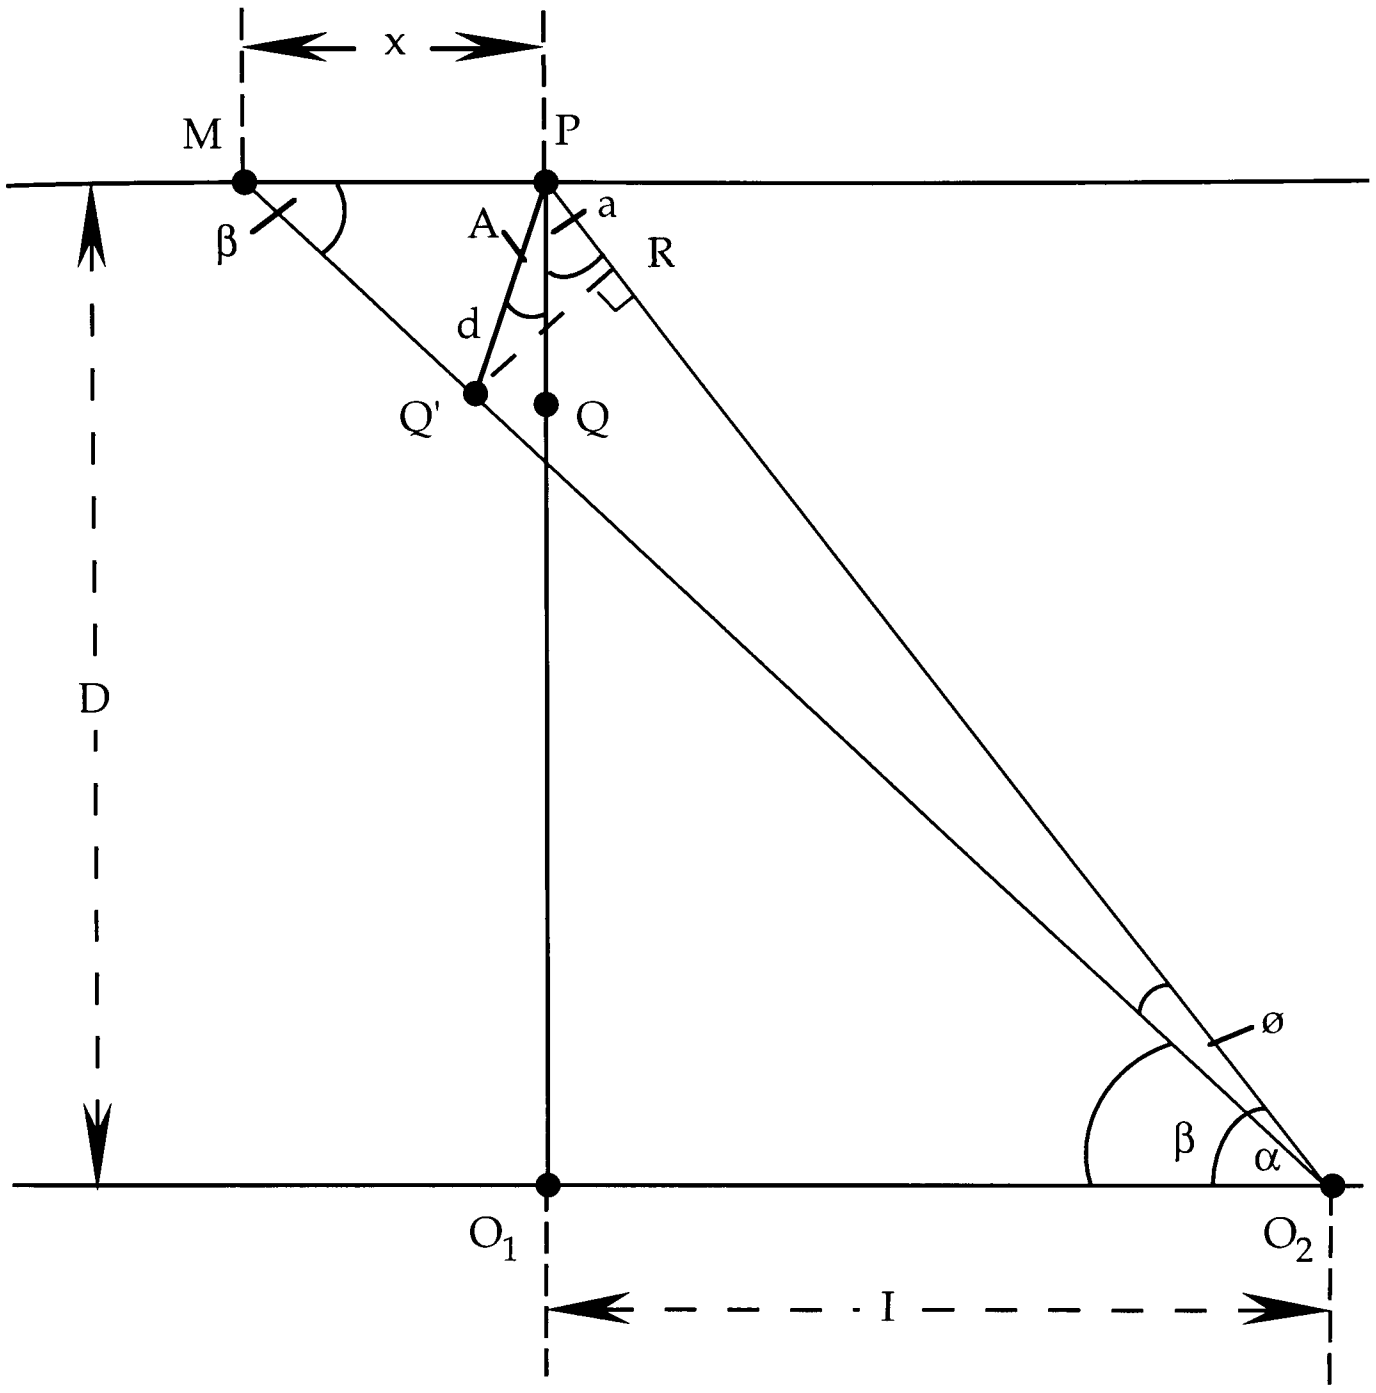


Fig. 5.5: Plan view of the motions of Points P and Q as the observer moves through distance I.

D	: viewing distance (from observer to axis of rotation at P)
d	: amplitude from rotation axis at P to peak at Q
I	: distance through which observer's eye moves
P, Q	: two points on the surface, P at the axis of rotation and Q (for example) at a peak
P, Q'	: the translated and rotated versions of P and Q
a	: angle through which the object rotates relative to the observer due to the translation of the observer from O_1 to O_2
A	: component of the object's rotation which is due to its own rotation about the vertical axis through P
x	: distance through which the dot situated on the peak of the surface translates as the observer's eye moves from O_1 to O_2

As the eye moves from O_1 to O_2, PQ undergoes a rotation of A.

$A = RPQ' - a$

$\sin RPQ' = RQ'/PQ' = RQ'/d$

$a = \tan^{-1}(I/D)$

$\therefore A = \sin^{-1}(RQ'/d) - \tan^{-1}(I/D)$.

Require RQ′.

Triangle $O_2Q'R$: $RQ'/O_2Q' = \sin ø$

$\therefore RQ' = O_2Q' \sin ø$.

But $ø = \alpha - \beta$

$$\begin{aligned}\therefore \sin ø &= \sin(\alpha - \beta)\\ &= \sin\alpha\cos\beta - \cos\alpha\sin\beta\\ &= [D/\sqrt{(D^2 + I^2)}][(I + x)/\sqrt{(D^2 + [I + x]^2)}] -\\ &\quad [I/\sqrt{(D^2 + I^2)}][D/\sqrt{(D^2 + [I + x]^2)}]\\ &= Dx/(\sqrt{(D^2 + I^2)}\sqrt{(D^2 + [I + x]^2)}).\end{aligned}$$

$O_2Q' = O_2M - MQ'$.

$O_2M = \sqrt{(D^2 + [I + x]^2)}$. Require MQ′.

Triangle MQ′P: $MQ'/\sin MPQ' = d/\sin Q'MP$.

But $MPQ' = 90° - A$ $\therefore \sin MPQ' = \cos A$

And $\sin Q'MP = \sin\beta = D/\sqrt{(D^2 + [I + x]^2)}$

$\therefore MQ' = \{d\cos A\sqrt{(D^2 + [I + x]^2)}\}/D$

$$\begin{aligned}\text{thus } OQ' &= \sqrt{(D^2 + [I + x]^2)} - \{d\cos A\sqrt{(D^2 + [I + x]^2)}\}/D\\ &= \{D\sqrt{(D^2 + [I + x]^2)} - d\cos A\sqrt{(D^2 + [I + x]^2)}\}/D\\ &= \{(D - d\cos A)\sqrt{(D^2 + [I + x]^2)}\}/D\end{aligned}$$

$$\begin{aligned}RQ' &= O_2Q' \sin ø\\ &= \{\{(D - d\cos A)/\sqrt{[D^2 + (I + x)^2]}\}/D\}\{Dx/\sqrt{(D^2 + I^2)}\sqrt{[D^2 + (I + x)^2]}\}\\ &= \{x\,(D - d\cos A)\}/\sqrt{(D^2 + I^2)}.\end{aligned}$$

$$\begin{aligned}\text{Thus, given that } A &= RPQ' - a\\ &= \sin^{-1}(RQ'/d) - \tan^{-1}(I/D),\end{aligned}$$

$A = \sin^{-1}\{(x[D - d\cos A])/(d\sqrt{[D^2 + I^2]})\} - \tan^{-1}(I/D)$.

In order to plot a family of theoretical curves for d and A, d was computed for values of A lying between ±25°, for three different relative motion amplitudes at a viewing distance, D, of 57 cm and an observer translation distance of 6.5 cm. Fig 5.6 below shows these curves for Peak-to-Trough amplitudes of 5, 10 and 20 arc min equivalent disparity.

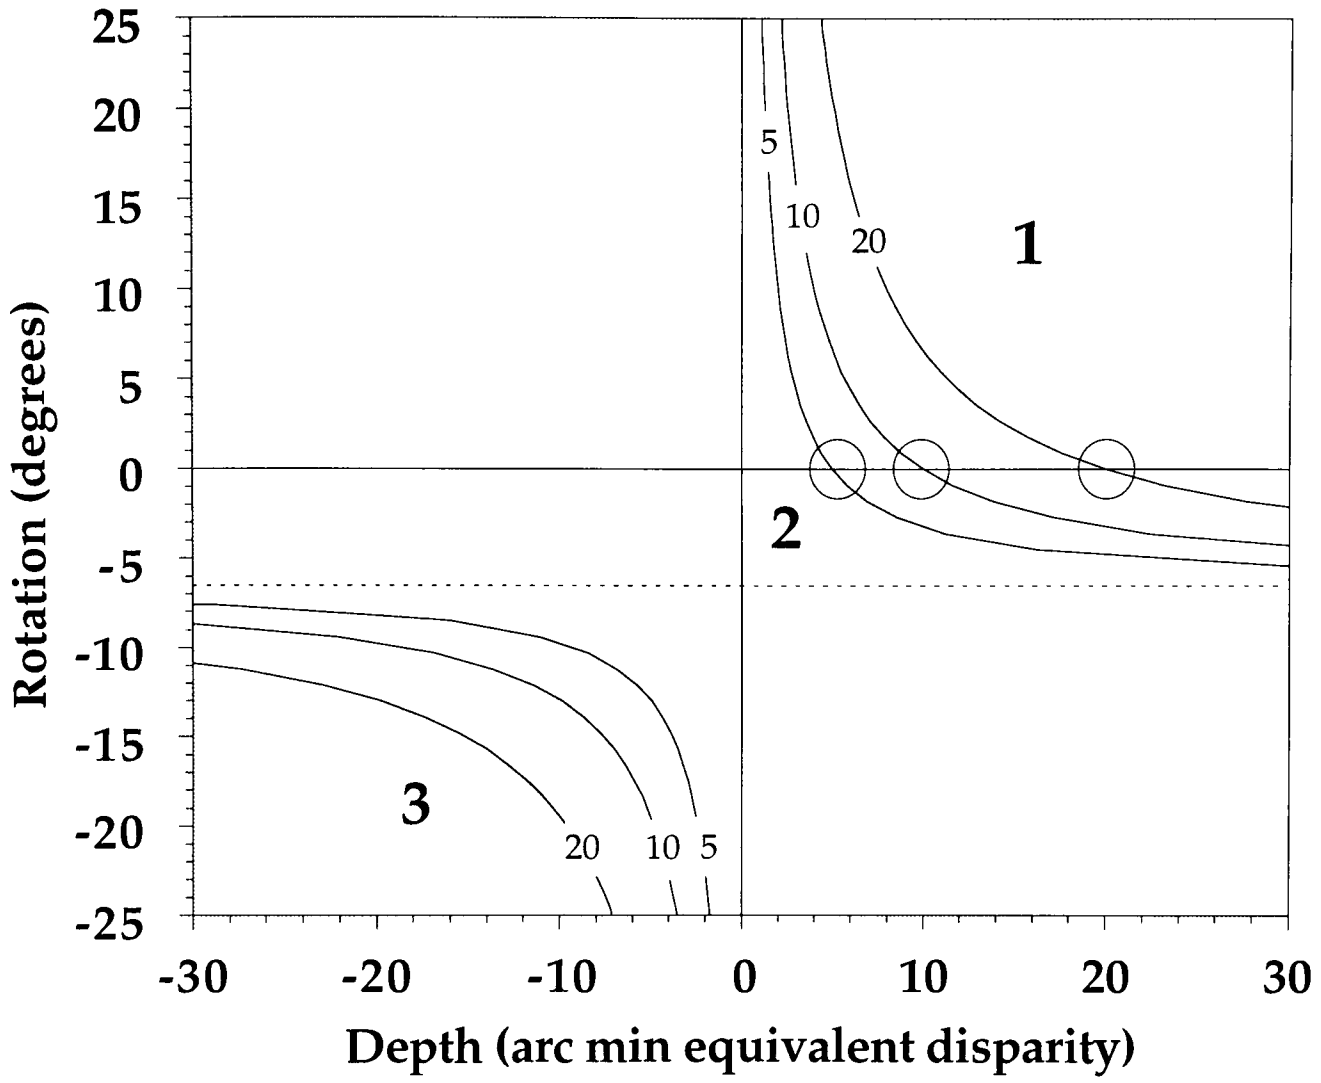


Fig 5.6: Depth and rotation combinations allowed by the geometric model, for Peak-to-Trough relative motion amplitudes of 5, 10 and 20 arc min equivalent disparity.

It can be seen that the solution set for Depth and Rotation for a given relative motion amplitude shifts position in [Depth, Rotation] space as the amplitude changes. The shape of the solution set also changes. The curves shift outwards as the amplitude increases, meaning that the range of allowed depths shifts to greater values also. The steepest portion of the curve, where the Rotation magnitude increases sharply, also moves to larger depth values.

Regions 1, 2 and 3 are marked on the graph.

Region 1 contains values of A which produce a solution where the surface rotates outwards in a 'convex' manner (see Figs. 1.4C and 5.3A).

The d values are *less* than those which would be predicted when A = 0° (i.e. when there is no rotation). The depth order is positive, i.e. Q' lies in front of P, relative to the observer.

Region 2 lies between the abscissa and the line A = -6.5°. This region contains values of A which produce a solution where the surface rotates inwards in a 'concave' manner (see Figs. 1.4B and 5.3B). The d values are greater than those which would be predicted when A = 0°, but the depth order is still positive. This portion of the graph arises because of the translation of the observer through 6.5 cm.

Region 3 lies beneath the line A = -6.5° and to the left of the line d = 0 arc min equivalent disparity. This region contains values of A which produce a solution where the surface rotates inwards in a 'concave' manner. The depth order is negative, i.e. Q' lies behind P, relative to the observer (see Fig. 1.4D). In Fig. 5.7, for example, Point Q' would be seen as lying behind Point P. Note that a negative depth order necessitates a concave rotation greater than 6.5°: a stationary surface of negative depth order is not permitted by the model.

The points at which each curve crosses the abscissa are circled. At each of these points, the surface is stationary (i.e. A = 0°) and the motion parallax is interpreted solely as depth. This case is geometrically analogous to stereoscopic disparity (see Figs. 1.2 and 1.3).

The horizontal dashed line represents the case where the relative motion amplitude is zero. The solution for zero parallax amplitude and non-zero depth is the case where the surface rotates in a concave manner, such that it always remains orthogonal to the observer's line of sight. The depth of the surface is indeterminate. The value of A which produces this result is that due to the observer translation. In the present case, A = -6.5°.

Experiments of this chapter also examine the perceived magnitudes of depth and shearing in square wave surfaces (Fig. 5.7). A model illustrating the geometric relationship between these attributes is presented in Fig. 5.8.

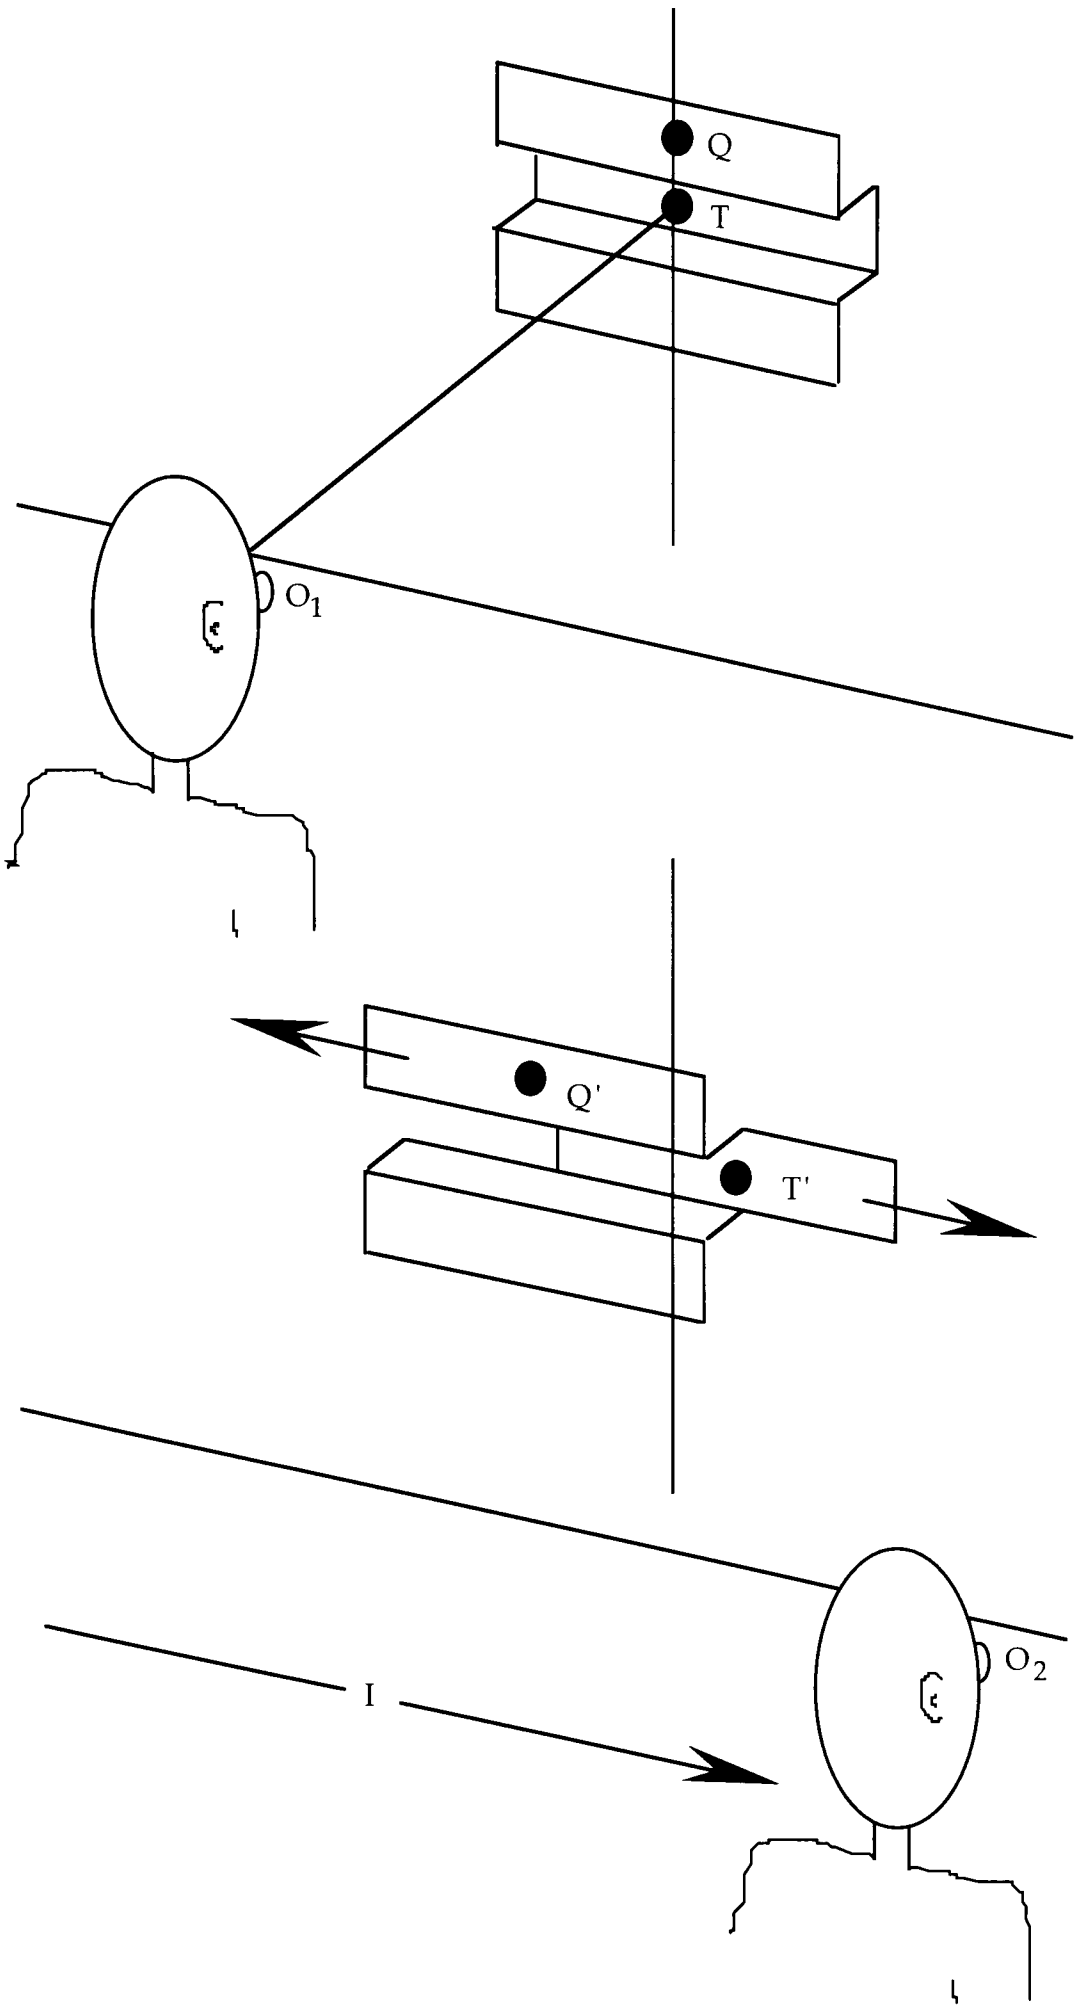


Fig 5.7: The nearer and further portions of a square wave surface translate in opposite directions ('shear') as the eye of the observer moves from O_1 to O_2 through distance I. Points Q and T (diagram 1) are situated on the nearer and further surfaces, respectively. Points Q′ and T′ (diagram 2) are the new positions of Points Q and T.

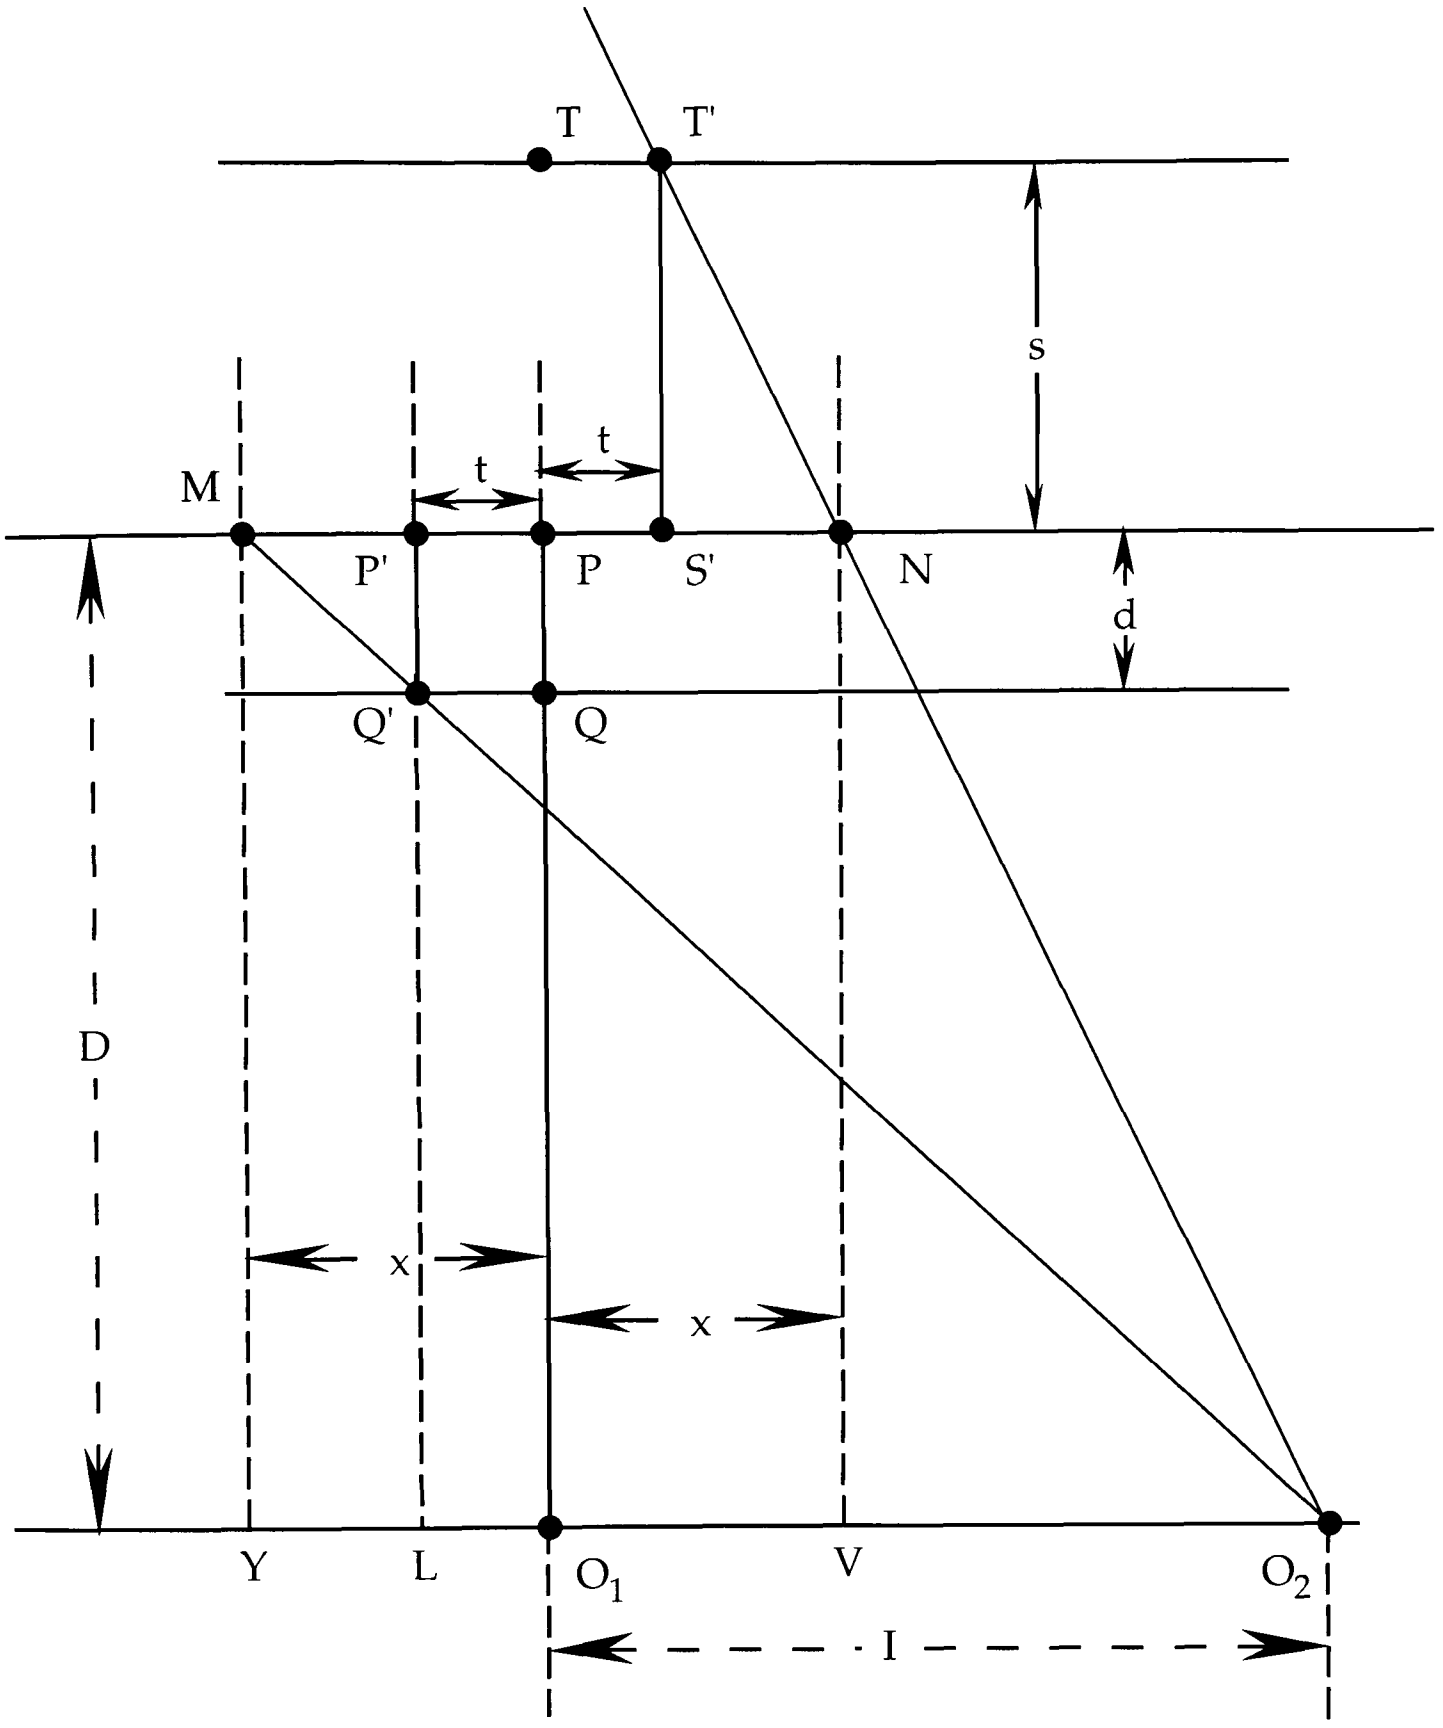


Fig 5.8: Plan view of the motions of Points Q and T as the observer moves through distance I.

D : viewing distance (from observer to monitor surface at P)

P, Q, T : three points, P at the monitor distance, and Q and T on the nearer and further surfaces, respectively, of the square wave

d : distance between monitor surface at P to nearer surface of the square wave at Q

s : distance between monitor surface at P to further surface of the square wave at T

I : distance through which observer's eye moves

x : distance through which the dot situated on the peak of the

surface translates as the observer's eye moves from O_1 to O_2

Q', T' : the translated versions of Q and T

t : distance through which points Q and T translate as the observer's eye moves from O_1 to O_2.

As the eye moves from O_1 to O_2, PQ undergoes a translation of magnitude t.
Similar triangles MYO_2 and $Q'LO_2$:

$$D/(I + x) = (D - d)/(I + t)$$
$$\therefore D\text{-}d = [D(I + t)]/(I + x)$$
$$d = D - [D(I + t)]/(I + x).$$

As the eye moves from O_1 to O_2, PT undergoes a translation of magnitude t.
Similar triangles T'S'N and NVO_2:

$$s/(x - t) = D/(I - x)$$
$$\therefore s = D(x - t)/(I - x).$$

$$\begin{aligned}\text{Total depth} &= d + s \\ &= \{D - [D(I + t)]/(I + x)\} + D(x - t)/(I - x) \\ &= [D(I + x)(I - x) - D(I + t)(I - x) + D(x - t)(I + x)]/[(I + x)(I - x)] \\ &= [D(I^2 - x^2 - I^2 + tx - tI + xI + xI - tx - tI + x^2)]/[(I + x)(I - x)] \\ &= D(2xI - 2tI)/[(I + x)(I - x)] \\ &= D\{2I(x - t)/[(I + x)(I - x)]\}.\end{aligned}$$

Fig 5.9 shows the combinations of depth and shearing permitted by the geometric model, for three values of relative motion at a viewing distance, D, of 57 cm, for an observer translation of 6.5 cm. Shearing values correspond to the distance P'S' (Fig. 5.8), expressed in equivalent disparity. It can be seen that the equations of the functions are of the form $y = k - x$, where k is the relative motion amplitude. This equation merely expresses the fact that the sum of the depth and shearing (expressed in arc min equivalent disparity) equals the relative motion amplitude.

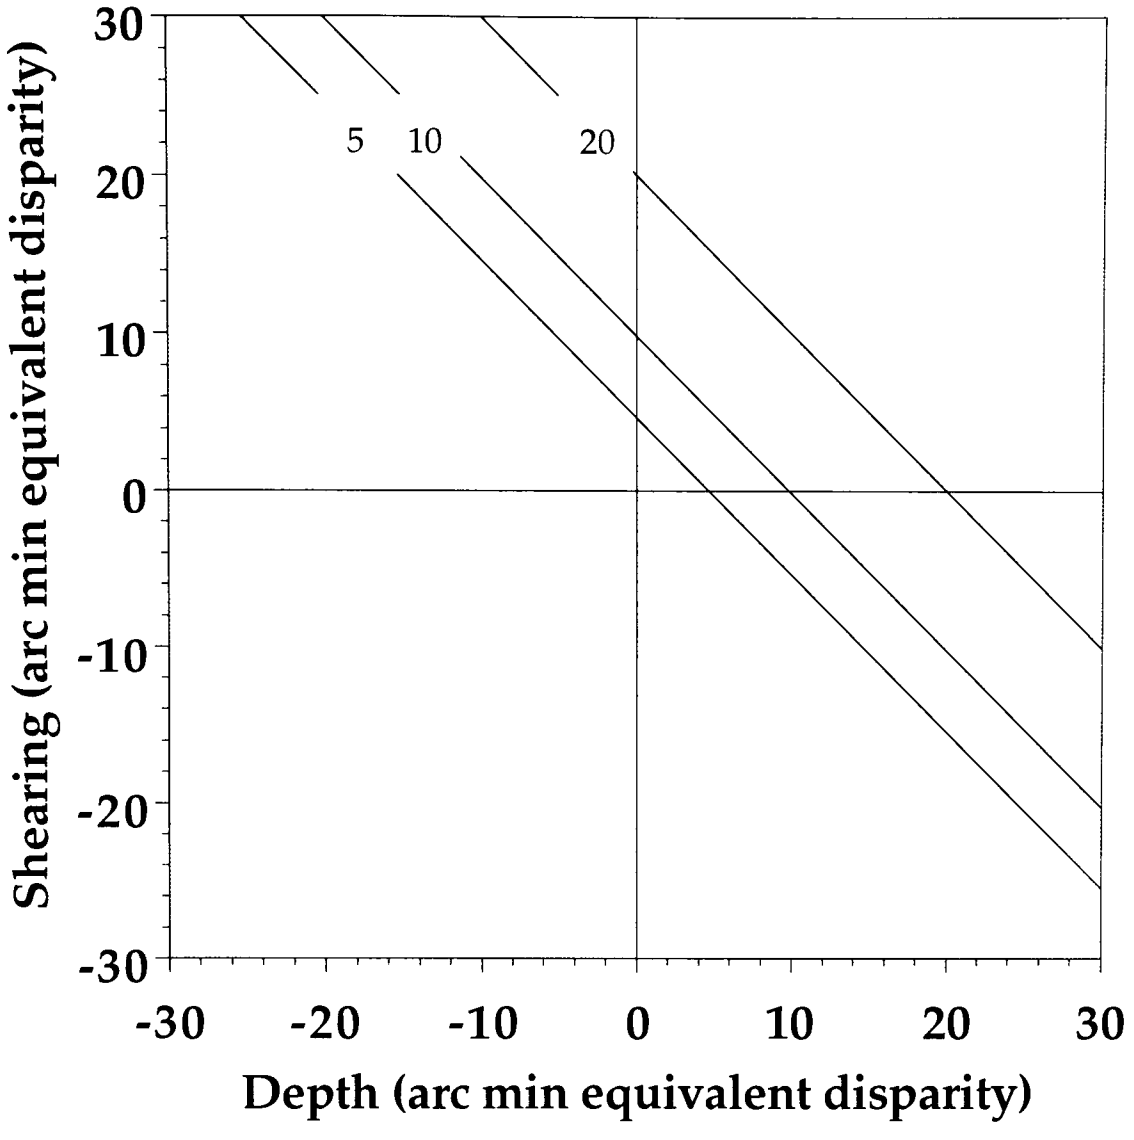


Fig. 5.9: Depth and shearing combinations allowed by the geometric model, for relative motion amplitudes of 5, 10 and 20 arc min equivalent disparity.

In order to make the relative magnitudes of the depth and shearing more intuitive, they are replotted below, expressed in mm (Fig. 5.10).

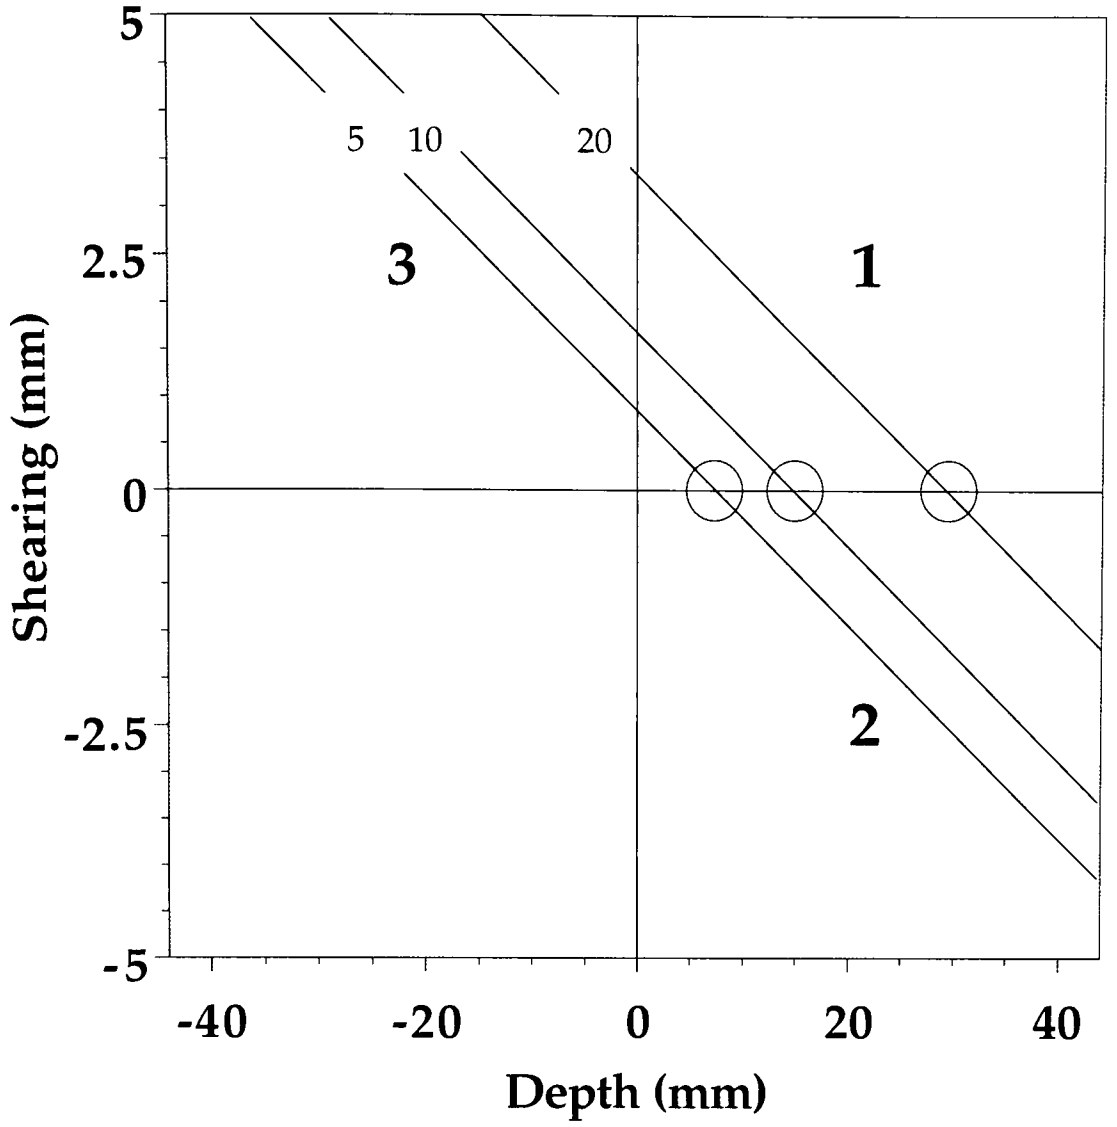


Fig. 5.10: Depth and shearing combinations allowed by the geometric model, for relative motion amplitudes of 5, 10 and 20 arc min equivalent disparity.

Regions 1, 2 and 3 are marked on the graph.

Region 1, the upper right quadrant, contains shearing values which lie between 0 and the extent of 2-D relative dot motion. The depth magnitudes are *less* than those which would be predicted when there is zero shearing. The depth order is positive.

Region 2, the lower right quadrant, contains negative shearing values. The depth magnitudes are greater than those which would be predicted when there is zero shearing. The depth order is positive.

Region 3, the upper left quadrant, contains shearing values which are greater than the extent of 2-D relative motion. The depth order is negative. In Fig. 5.8, for example, Point Q′ would be seen as lying behind Point T′.

The points at which each curve crosses the abscissa are circled. At each of these points the surface is stationary — there is no shearing — and the

motion parallax is interpreted solely as depth.

Note that neither model incorporates the effects of perspective changes or the acceleration component of the flow field which accompany a particular amount of rotation. Although the visual system is reasonably accurate in tasks involving velocity extraction (McKee, 1981), the detection of acceleration has been shown to be relatively poor (Snowden and Braddick, 1991; Werkhoven, Snippe and Toet, 1992; Snippe and Werkhoven, 1993). It has been suggested that the main descriptor of human structure-from-motion performance does not involve acceleration information (Van Veen and Werkhoven, 1996).

The influence of perspective information on perceived rotation magnitude is considered in Chapter 6 (see also Experiments 3, 9 and 12).

The relationship between the relative magnitudes of perceived depth and perceived motion is investigated using two types of surface. Experiments 11, 12, and 15 use sinusoidally corrugated surfaces of fixed peak-to-trough amplitude, and observers make depth and rotation matches. Experiments 13 and 14 use square wave surfaces, also of fixed amplitude, and observers make depth and translation/shearing matches. The perceived amounts of depth and motion in each case are compared with the values predicted by the relevant geometric model.

There are at least three hypotheses concerning the nature of the co-variation of perceived depth and motion. Firstly, the co-variation may not exist at all: manipulation of a given stimulus attribute (for example size or corrugation frequency) may lead to variations in the perceived magnitudes of depth and motion which are in the same direction: both magnitudes might increase, or both might decrease. It has previously been suggested that the perception of depth and rotation are governed by independent processes (Braunstein, 1976) and, as outlined above, evidence suggests that under certain conditions there is no trade-off relationship (Liter et al., 1992; 1993).

Secondly, the co-variation may be exactly in line with the predictions of the geometric models. That is, for a given amplitude of relative motion the relative magnitudes of perceived depth and rotation (or shearing) may fall on or near to the relevant line of permissible solutions. Williams (1993) and Van Veen and Werkhoven (1996) found this to be the case for a range of rotating surfaces.

Finally, there may exist a degree of co-variation such that an increase in the magnitude of perceived motion goes hand in hand with a decrease in the magnitude of perceived depth (or vice versa), but the trade-off may not be in quantitative agreement with the model.

The aim of the first two experiments of this chapter is to examine the relationship between the relative magnitudes of perceived rotation and depth. Manipulation of both corrugation frequency and size lead to variations in the magnitude of perceived rotation (see Experiments 7a, 7b and 9). Experiments 11 and 12 monitor the concomitant variations in perceived depth.

5.2 Experiment 11: Matched Rotation and Matched Depth as a function of Corrugation Frequency

5.2.1 Stimuli

Stimuli contained a corrugation frequency of 0.05, 0.1. 0.2 or 0.4 cpd. Stimuli were 20° in diameter. Five equally spaced values of relative motion were used: 4, 12, 20, 28 and 36 arc min equivalent disparity. The peak-to-trough depths simulated by these extents of relative motion are 6, 18, 29, 41 and 53 mm respectively.

5.2.2 Procedure

Each trial required the subject to make a Depth Match and a Rotation Match. Trials were blocked in groups of ten, corresponding to two repetitions of each of the relative motion amplitudes. Subjects completed three sessions of one block per corrugation frequency per session. The order of blocks within each session was randomised. In half the total

number of trials the depth judgement was made first and the rotation judgement was made second. In the other half, this order was reversed. Rotation Matches were obtained using the method described in Chapter 4. Depth Matches were obtained by asking subjects to adjust a pair of digital callipers to match the amount of perceived peak-to-trough depth in the stimulus.

5.2.3 Results

Figs. 5.11 and 5.13 show Matched Rotation as a function of relative motion for the four corrugation frequencies, for Subjects SAS and SJP respectively. Figs. 5.12 and 5.14 show Matched Depth as a function of relative motion for the four corrugation frequencies, for Subjects SAS and SJP respectively. The dashed line on each of these graphs represents simulated depth. The graphs show that as corrugation frequency increased, Matched Rotation increased and Matched Depth decreased. This effect was especially noticeable for the three lowest corrugation frequencies; data for the 0.4 cpd stimulus were not particularly well separated from those of the 0.2 cpd stimulus.

Figs. 5.15 and 5.16 replot the data to show its relationship with the depth/rotation combinations predicted by the model for each amplitude of relative motion. The left-hand legend shows the relative motion amplitude (arc min equivalent disparity) which produces each theoretical depth/rotation curve. The right-hand legend shows the surface corrugation frequency. Symbols of identical colour arose from surfaces containing identical relative motion amplitudes, the values of which are indicated in the former legend. The data should theoretically lie on the curve of the same colour. It is evident that, for all surfaces tested, there was an over-estimation of the combined amounts of depth and rotation relative to those predicted by the geometric model.

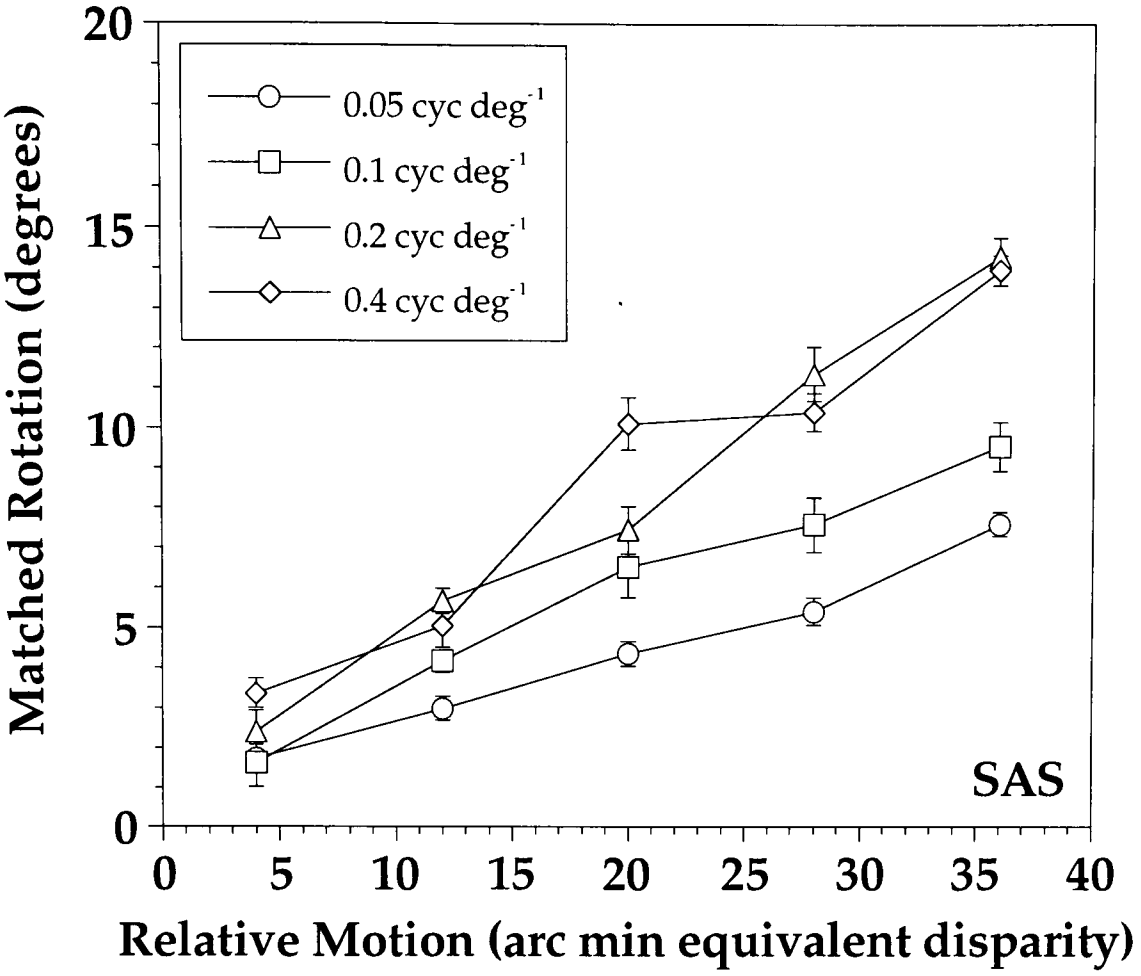


Fig. 5.11: Matched Rotation as a function of Relative Motion for four corrugation frequencies, for Subject SAS.

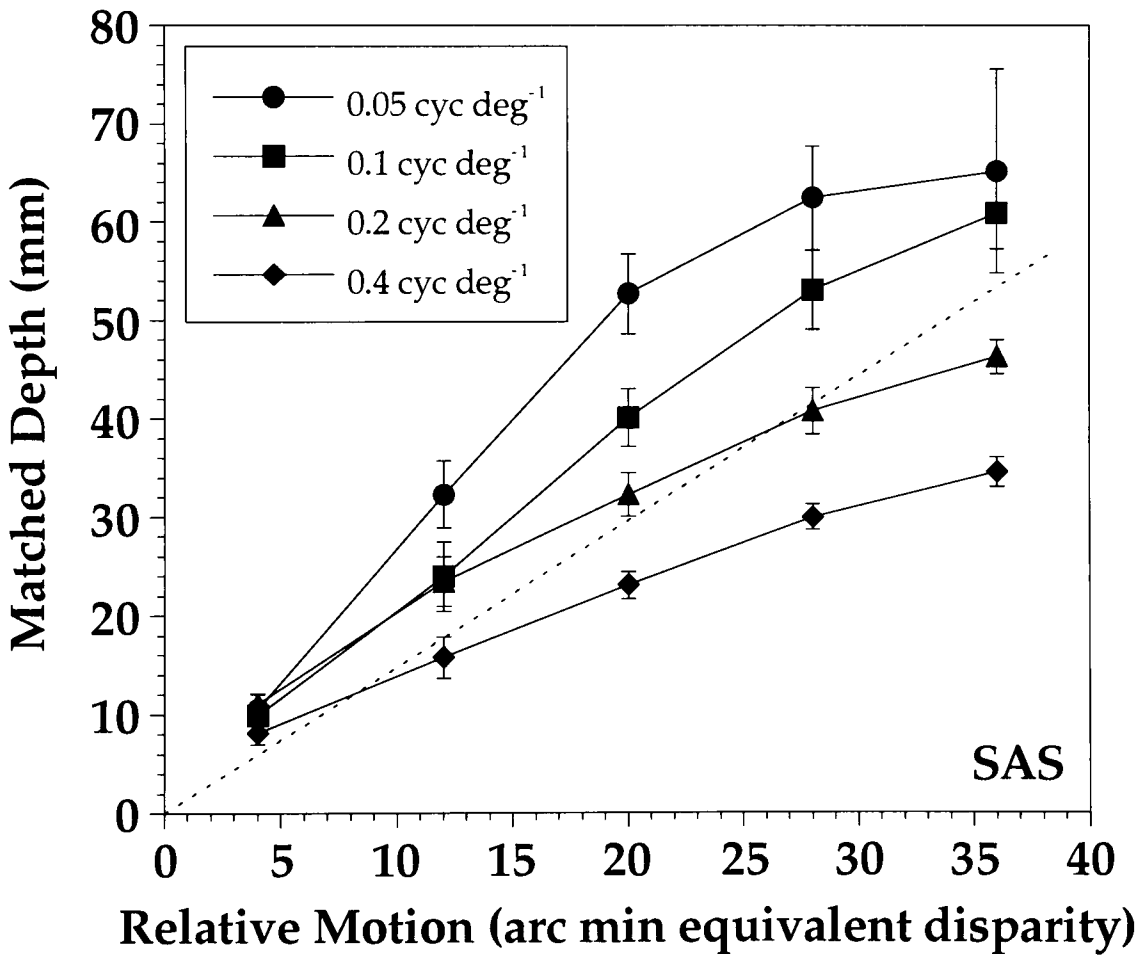


Fig. 5.12: Matched Depth as a function of Relative Motion for four corrugation frequencies, for Subject SAS. The dashed line represents simulated depth.

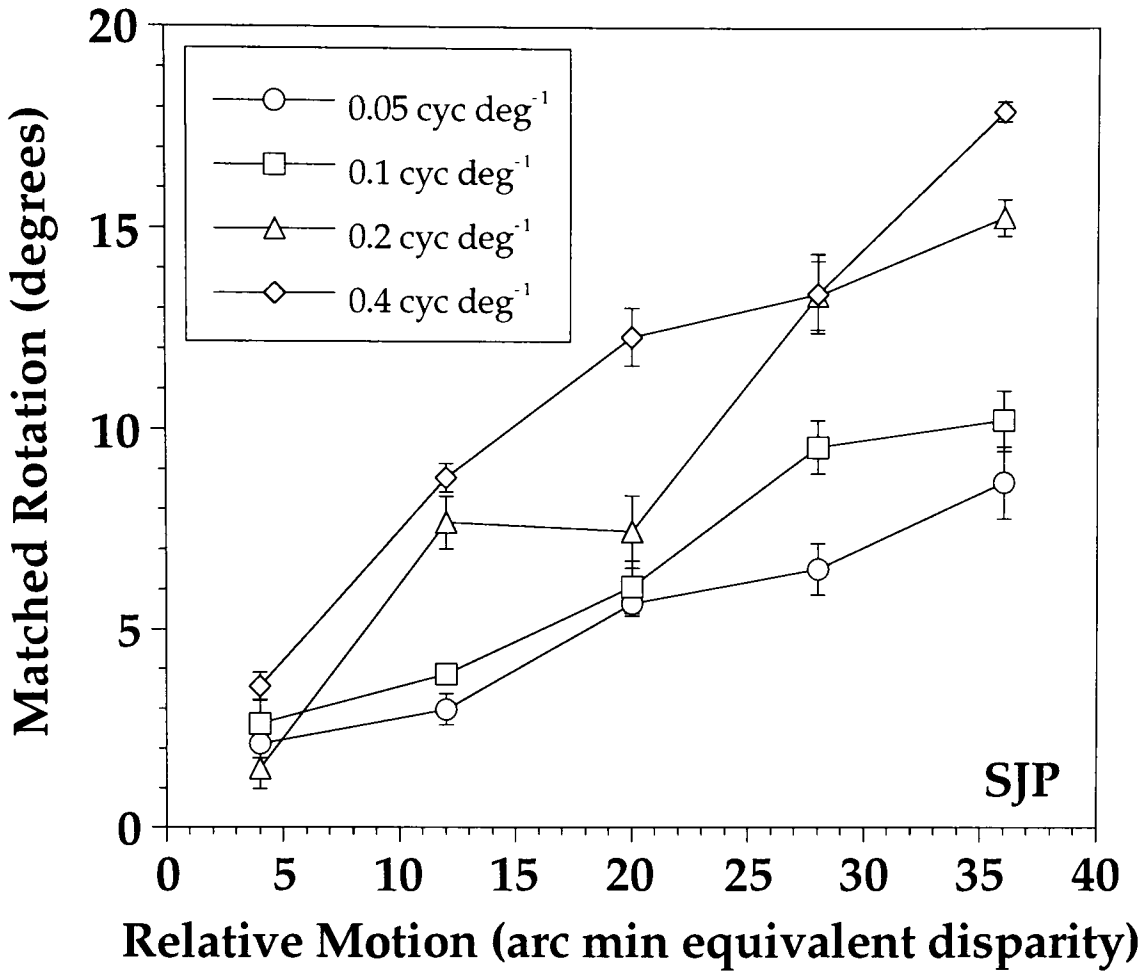


Fig. 5.13: Matched Rotation as a function of Relative Motion for four corrugation frequencies, for Subject SJP.

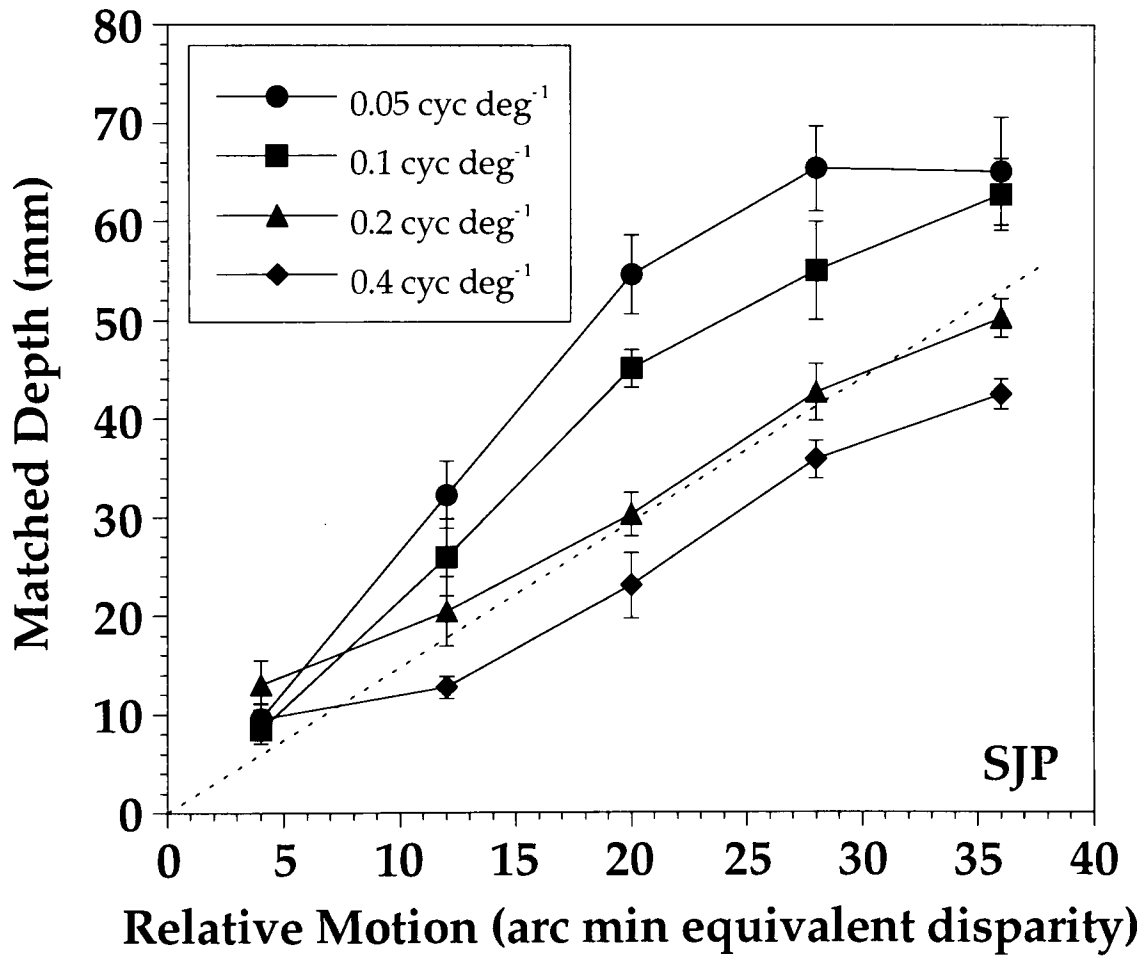
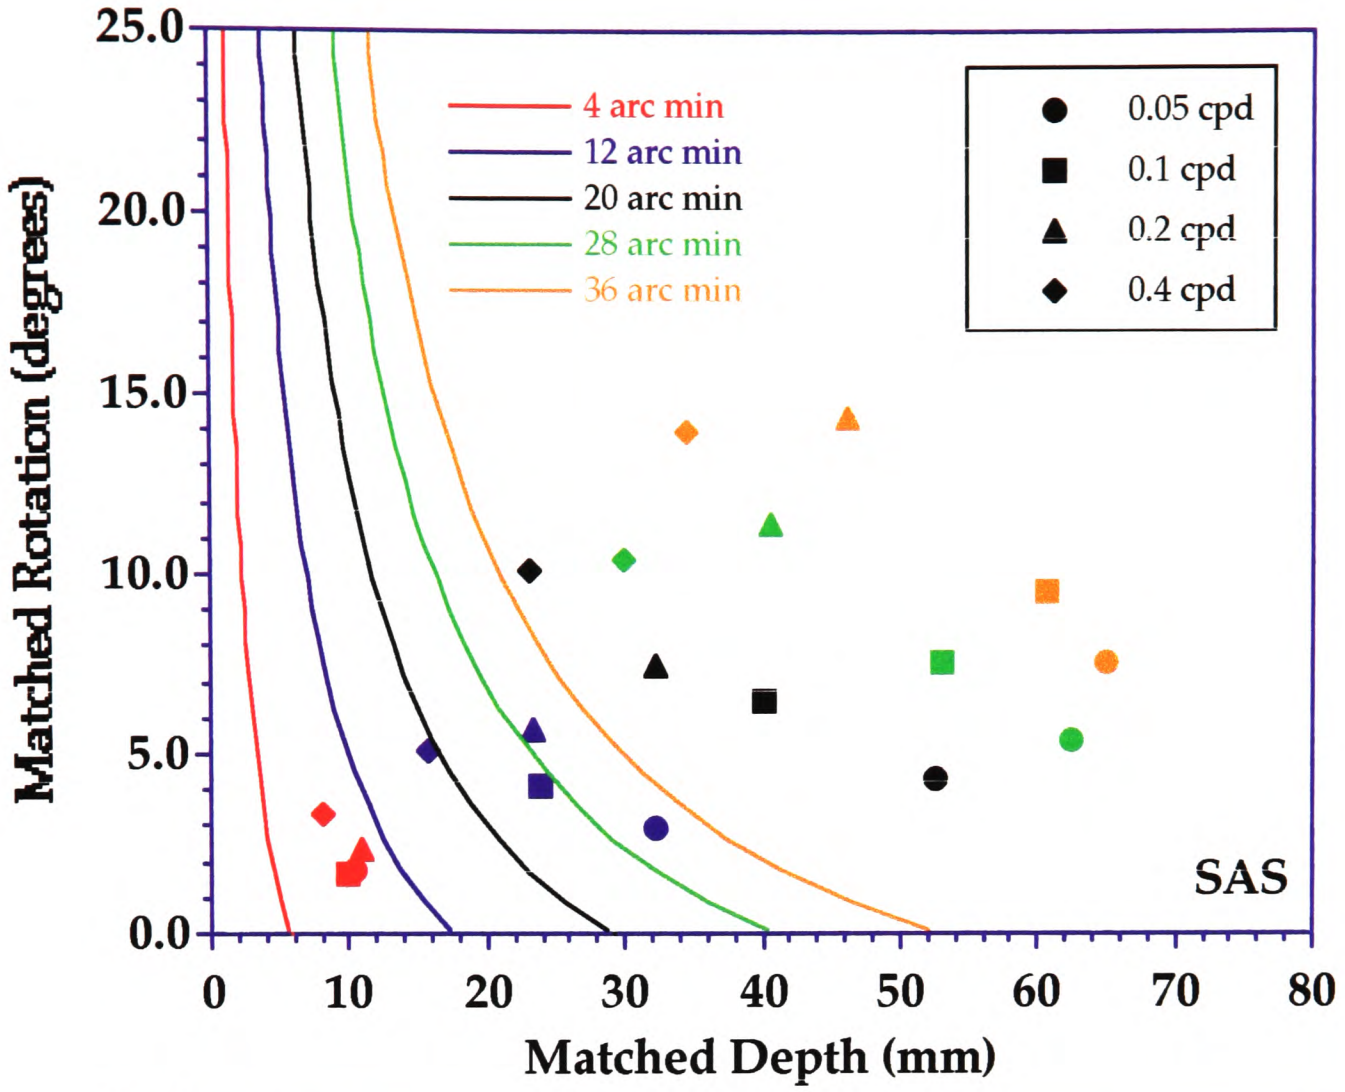
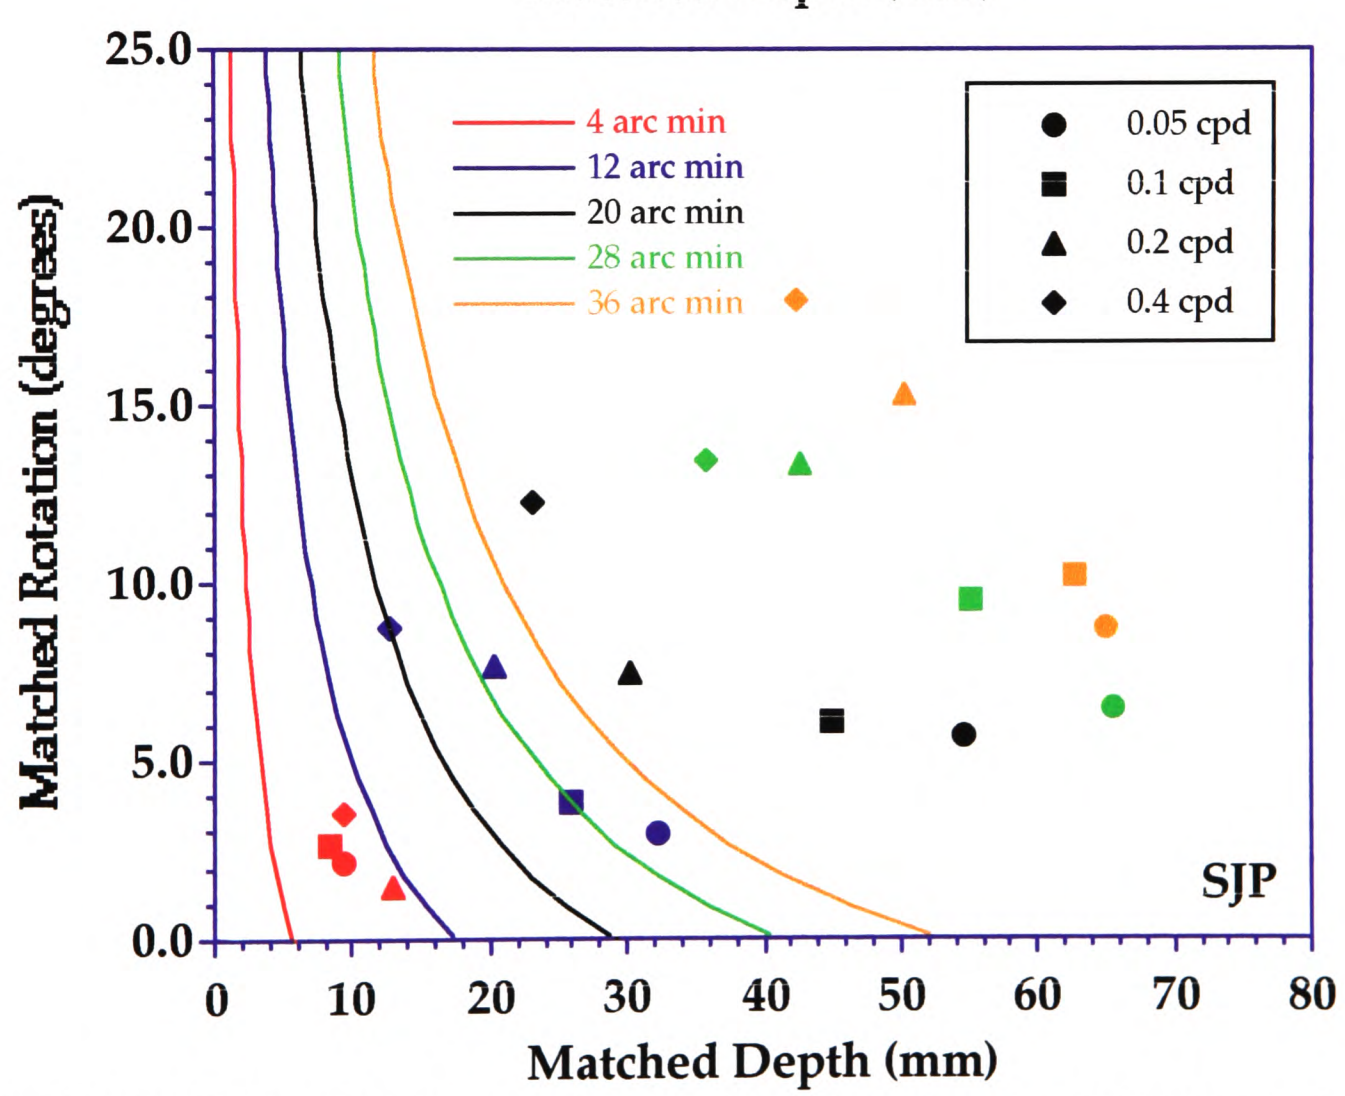


Fig. 5.14: Matched Depth as a function of Relative Motion for four corrugation frequencies, for Subject SJP. The dashed line represents simulated depth.





Figs. 5.15 & 5.16: Matched Rotation versus Matched Depth for four corrugation frequencies, for Subjects SAS and SJP. The curves show the rotation/depth combinations predicted by the geometric model for the five different relative motion amplitudes. Data should theoretically fall on the line of the same colour.

Figs. 5.17 and 5.18 quantify the extent of over-estimation of the combined depth/rotation values for each surface. The figures show the relative motion values which would theoretically be required to produce the perceived amounts of rotation and depth in each surface. Thus, for example, Subject SAS perceived a 0.1 cpd with a relative motion amplitude of 20 arc min equivalent disparity as having a peak-to-trough depth of approximately 40 mm, and to rotate through approximately 6.5° (Fig. 5.15). A relative motion amplitude of approximately 55 arc min equivalent disparity is geometrically consistent with these magnitudes (Fig. 5.17). It can be seen that the relative motion amplitudes theoretically required to produce the observed magnitudes of depth and rotation are significantly higher than the physical amplitudes of the surfaces. The predicted amplitudes are generally between twice and three times the physical amplitudes, and are larger for Subject SJP than for Subject SAS. At low relative motion amplitudes there is little difference between the predicted amplitudes of each surface. As the amplitude increases, the discrepancy between the predicted and physical amplitudes decreases slightly with increasing corrugation frequency.

5.2.4 Discussion

The finding that perceived rotation increases with increasing corrugation frequency has previously been noted and discussed in Section 4.3. That this increase goes hand in hand with a decrease in perceived depth suggests there is a trade-off between these two stimulus attributes: depth is perceived at the expense of rotation and vice versa. Figs. 5.15 and 5.16 show that in almost all cases, the combined magnitudes of perceived depth and rotation are larger than predicted by the geometric model. One interpretation of these data is that observers are *not* quantitatively constrained by the geometric model in their interpretation of a given amplitude of relative motion.

Figs. 5.17 and 5.18 quantify the relationship between the physical amplitude of each surface and the amplitude of relative motion theoretically required to produce the observed magnitudes of perceived depth and rotation of that surface. The fact that the discrepancy between the two amplitudes generally decreases with increasing corrugation frequency, at least at the higher relative motion amplitudes, could be due

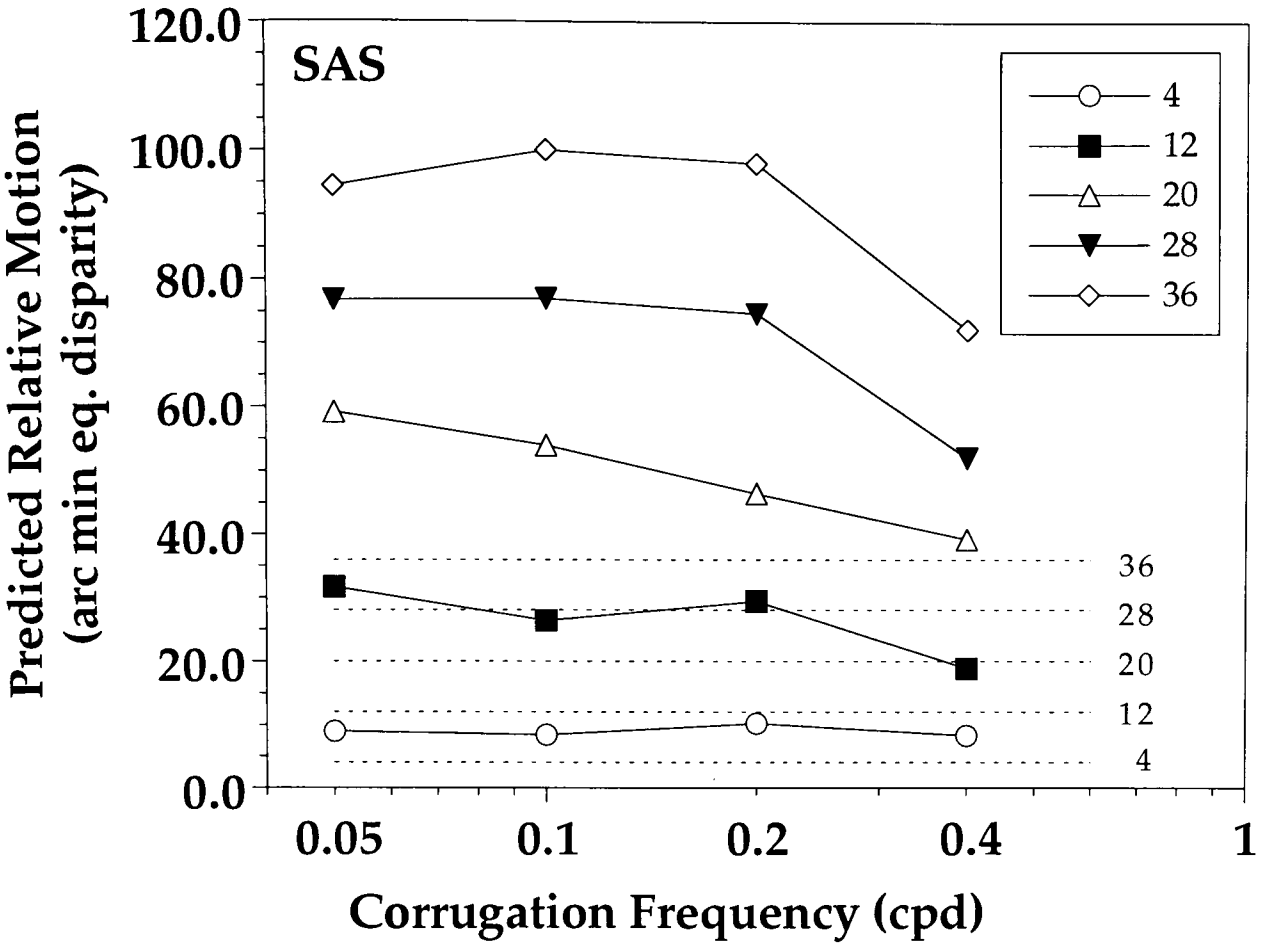


Fig 5.17: The relative motion values which would theoretically be required to produce the perceived amounts of rotation and depth in each surface, for Subject SAS. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

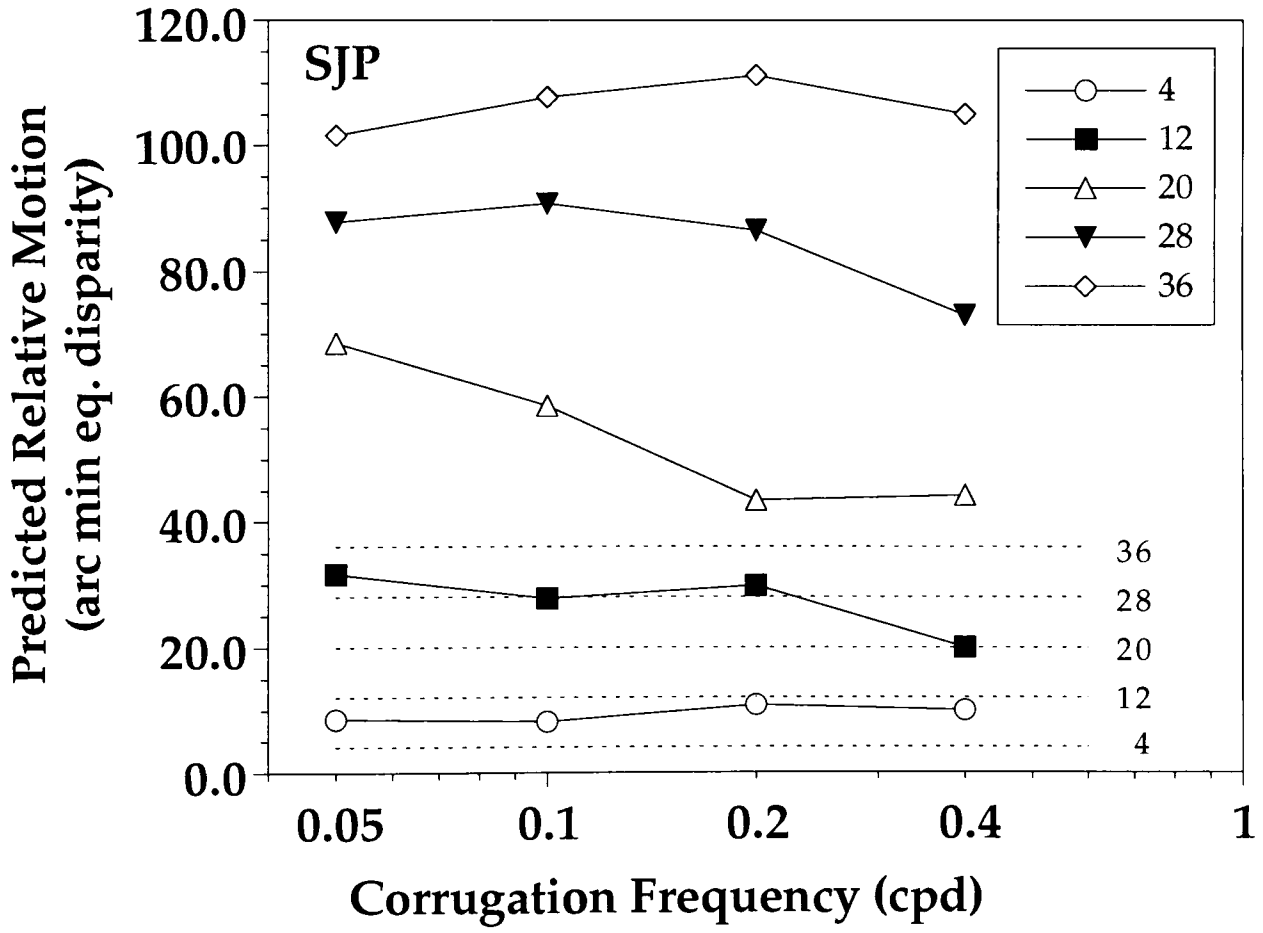


Fig 5.18: The relative motion values which would theoretically be required to produce the perceived amounts of rotation and depth in each surface, for Subject SJP. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

to the shape of the depth/rotation curves of the model. The increases in perceived rotation magnitude which occur with increasing corrugation frequency are consistent with relatively small increases in the theoretical relative motion amplitude, whereas the increases in perceived depth which occur with decreasing corrugation frequency are consistent with relatively large increases in the theoretical amplitude. The discrepancy between the physical and theoretical amplitudes of the 0.4 cpd surface is noticeably smaller than the amplitude discrepancy of the other surfaces. This may partly be due to a proportion of the relative motion being interpreted as surface non-rigidity.

There are three possible explanations of the results. Firstly, it may be the case that relative motion is interpreted both as depth *and* as rotation — i.e. the motion is 'counted' twice. Obviously this explanation cannot fully account for the results since, with the exception of the 4 arc min equivalent disparity surfaces, the theoretical amplitudes which account for the observed data are frequently greater than twice the physical amplitudes. For surfaces containing 36 arc min equivalent disparity, the former amplitudes are generally nearer three times the latter.

The second possible explanation of the results is an over-estimation of the relative motion amplitude. In this case it would be unnecessary to posit that relative motion is 'counted' twice — it may merely be the case that a significantly larger amount of relative motion is perceived than is actually present and thus more motion is available for interpretation either as depth or as rotation.

Thirdly, the observed over-estimation may be due to the psychophysical method used to estimate the perceived magnitudes of depth and rotation. As noted in Section 2.4.3 the method of adjustment is not bias-free and thus it is entirely possible that observers systematically over-represented their perceived magnitudes of depth and rotation. The likelihood of any of these three explanations is discussed further in Section 5.4.

Rogers and Collett (1989) presented observers with surfaces defined by both motion parallax and binocular disparity, and found that the perceived depth magnitude was determined largely by the latter cue. The

former cue was incorporated into the percept as a rotation of the structure, as described in Section 1.6.1. When peak-to-trough disparity was held constant at 2 arc min, and the relative motion amplitude was varied between 2 and 12 arc min equivalent disparity, it was found that a greater amount of depth was perceived as the surface corrugation frequency decreased. The authors suggested that this was due to the decreasing influence of the disparity cue on perceived depth as the corrugation frequency (and thus the disparity gradient) decreased. The authors additionally observed (albeit non-quantitatively) that the amount of rotation also increased with decreasing corrugation frequency. Their explanation of this is as follows. When disparity gradients are shallow — as in the case of surfaces of low corrugation frequency — a discrepancy between the amounts of parallax and disparity generates a relatively large rotational discrepancy. For a surface of high corrugation frequency, the same absolute discrepancy will generate a relatively small rotational discrepancy. In other words the degree of rotation required to 'explain' a given absolute discrepancy between the parallax and disparity amplitudes increases with decreasing corrugation frequency. The fact that perceived depth and perceived rotation both increase with decreasing corrugation frequency obviously conflicts with predictions based upon a depth-rotation trade-off. However, the central tenet of Rogers and Collett's (1989) paper is that the visual system perceives a structure whose depth and rotational characteristics minimise the discrepant disparity and parallax signals. Under this assumption it is unsurprising that the depth/rotation relationship is not strictly adhered to, since this would take account of the motion parallax information only.

The effect of disparity gradient on perceived depth magnitude

Considering the depth data alone, the finding that less depth is perceived with increasing disparity gradient has been demonstrated using stereoscopic stimuli. Bülthoff, Fahle and Wegmann (1991) found that the perceived depth between two stimuli decreased with increasing disparity gradient when absolute disparity remained constant. Observers were presented with stimuli which contained two points in both eyes, separated either horizontally, vertically or diagonally. One point had crossed disparity, the other point had uncrossed disparity; the stimulus thus appeared to contain a point on a plane beyond fixation and a point on a

plane in front of fixation. The task was to decide which of a set of reference lines, separated in depth, was closest to the nearer point and which was closest to the further point. For a given disparity value, various disparity gradients were created by moving the points towards or away from each other along paths within fronto-parallel planes. Disparities of 0 - 54 arc min were used. Disparity gradients ranged between 0.3 and 1.9.

It was found that for a given disparity, an increase in the disparity gradient led to a decrease in perceived depth separation between the two adjacent stimuli. This effect was especially marked when the stimuli consisted of lines. It was less well marked when two isolated points or two small symbols were used.

The authors suggested that the effect might be attributable to a tendency of the brain to make conservative depth estimates when faced with a stimulus containing line elements whose relative positions cannot be determined sufficiently accurately for an exact localisation in depth (Yuille, Bülthoff and Fahle, 1987). A spatial localisation uncertainty of a given size corresponds to a larger possible range of depths in steep gradient stimuli than in shallow gradient stimuli. The visual system chooses the interpretation with the smallest overall disparity difference in order to compensate for the potentially large depth error in these surfaces. Grosso, Sandini and Tistarelli (1989) have suggested a computational approach whereby signals are weighted according to their 'robustness' or reliability, and the down-scaling of depth in steep gradient stimuli might reflect the operation of this general principle.

In keeping with the results of the present experiment, the data of Bülthoff at al. (1991) show that under certain circumstances less depth is perceived in high frequency than in low frequency stimuli. A number of differences exist between the stimuli of that study and the stimuli of this thesis, however, so application of their results to the current experiment may be rather limited. Bülthoff et al. (1991) used displays which subtended only about 1° whilst the current experiment used displays of diameter 20°. Moreover, the displays of Bülthoff et al. consisted of single points or lines, whereas the displays of the current experiment simulated continuous, 3-D surfaces for which the disparity gradient scaling effect may be less strong.

A second study which partly bears on the findings of Experiment 12 was carried out by Hirsch and Weymouth (1948a, 1948b). They found that stereoacuity decreased as the distance between a comparison and a test element increased: a greater amount of disparity was required to reach the threshold. From these threshold data Howard and Rogers (1995) predicted that at supra-threshold disparities the results of Bülthoff et al. (1991) might be reversed if their experiment was repeated using shallower disparity gradients. A reversal of the trend would obviously also be reversed with respect to the pattern of results of the current experiment. As previously noted, however, the stimuli in the two situations are obviously highly dissimilar and it is not unlikely that different factors affect each set of results.

The 'Compactness Assumption'

Caudek and Proffitt (1993) noted that observers viewing SKE and motion parallax displays made depth magnitude judgements which were consistent with a 'compactness assumption'. Subjects viewed a simulated dihedral angle, or wedge, consisting of two slanted planes meeting at a horizontal ridge. In the motion parallax conditions, the motion gradient contained both object- and observer-relative transformations. The SKE conditions contained object-relative transformations only and thus did not contain sufficient information for the extraction of depth magnitude. For various wedge widths and depths, observers instructed the experimenter to adjust an icon in order that its size matched the perceived depth magnitude of the structure. It was found that an increasing amount of depth was perceived as the width of the wedge increased. When perceived depth (averaged across conditions) was plotted as a function of simulated depth, with both depths expressed as a proportion of the stimulus width, it was found that perceived depth magnitudes were never substantially greater than the object widths, consistent with the notion that observers' depth estimates were scaled by the projected stimulus size and did not generally exceed unity in terms of the depth-to-width ratio. This supported an earlier finding by Loomis and Eby (1988), who demonstrated that the perceived shape of rotating objects was influenced by the sizes of their projected outlines.

Although the perceived depths (expressed as a proportion of peak-to-

peak distance) of 16 of the 24 surfaces in the current experiment did not exceed unity and are thus consistent with the compactness assumption, these data do not directly speak to the validity of the assumption. The basic tenet of the assumption is that the perceived depths (expressed as a proportion of projected width) *of stimuli whose simulated depths* (expressed as a proportion of projected width) *exceed unity* do *not* exceed unity. In the current experiment, the only structures whose simulated depths (expressed as a proportion of peak-to-peak distance) exceed unity are the 0.4 cpd surface of amplitudes 20, 28 and 36 arc min equivalent disparity and the 0.2 cpd surface of 36 arc min equivalent disparity. Of these, the perceived depths of the 0.2 surface and the 0.4 cpd surface with amplitude 20 arc min equivalent disparity did not exceed unity for either subject. The pattern of the current data is also in accord with Caudek and Proffitt's (1993) finding that the perceived depth of a surface is scaled according to the width of the structure.

Durgin et al. (1995) also found evidence that the perceived depth of an object was affected by its projected width. They presented observers with simulated cones defined through either binocular disparity, motion parallax or a small simulated rotation. The base of the cone was 6 cm and the depth ranged from 40% to 300% of this value. Observers indicated the perceived depth magnitude of the cone using a mouse-adjustable icon. It was found that the depth in the motion parallax and object rotation displays was approximately 100% of the diameter of the base for all simulated depths. These results were interpreted as supporting the compactness assumption. The authors suggested that despite the fact that the cones in the motion parallax condition were simulated as translating, ".. they were, for practical purposes, indistinguishable from simulations of different-sized cones translating and rotating simultaneously". In this situation it is hypothesised that the visual system uses a perceptual heuristic such as the compactness assumption to estimate the objects' depths.

A heuristic as strong as the compactness assumption was not used in the current experiment. This is evident from Figs. 5.12 and 5.14, which illustrate that variations in the relative motion amplitude (analogous to Durgin et al.'s 'simulated depth') led to variations in perceived depth. Such variations would obviously not be expected under the compactness

assumption, since this states that the best guess of the system is that ".. the z-axis extent of the visual object equals the smallest of the dimensions exhibited by the projection of the object in the x-y plane", and thus predicts no effect of simulated depth. In this respect the data of Caudek and Proffitt themselves do not fully support the assumption, since variations in the simulated depth of their stimuli also resulted in variations in perceived depth.

Caudek and Proffitt's (1993) notion that the visual system makes a 'compactness assumption', and their observation that perceived depth scales with projected width, do not provide any information concerning the underlying basis for these observations. However, if there exists a depth/rotation trade-off relationship such that one attribute is perceived at the expense of the other, one possible explanation of Caudek and Proffitt's results is that less depth was perceived in the smaller stimuli because they were perceived to rotate more. Since the magnitudes of simulated wedge depth remained constant while the dimensions of the wedge were systematically varied, consequent variations in the motion gradients of the stimuli obviously occurred. Smaller stimuli would have had greater motion gradients. On the basis of Experiments 6 and 7 and the current experiment, the perceived rotation of these stimuli is likely to have been greater than that of larger stimuli. Geometrical considerations would thus predict that the perceived depth would increase with increasing wedge width — exactly the pattern of results observed.

Depth over-estimation

The depths simulated by each value of relative motion (assuming that none of the motion is interpreted as rotation) are represented in Figs. 5.15 and 5.16 by the points at which the curves cross the abscissa. It can be seen that observers significantly over-estimated the depth in the 0.05 and 0.1 cpd surfaces. This is especially surprising in the light of the fact that convex rotation was also perceived in all surfaces, since the perception of such rotation should theoretically reduce the amount of perceived depth. The extent of the depth over-estimation was greatest in the 0.05 cpd condition and least in the 0.4 cpd condition. In the latter condition, matched depth was under-estimated at the higher relative motion values.

A number of other studies involving depth matching have also found that over-estimation of depth occurs when relative motion amplitudes are small. Ono, Rivest and Ono (1986) presented observers with motion parallax surfaces and asked the observers to indicate the amount of perceived depth by adjusting the distance between two rods. At a viewing distance of 40 cm the amount of perceived depth was over-estimated in each of two differently-sized displays. The size and corrugation frequency of the large display were 40° x 42° and 0.04 cpd respectively. The size and corrugation frequency of the smaller display were respectively half and twice these values. The equivalent disparity was 0.47° for both displays. This corresponds to a simulated depth of 2.3 cm. The mean apparent depth of the displays was over-estimated. The mean perceived depths of the small and large displays were 2.6 and 4.0 cm respectively.

In a second experiment observers were asked to adjust the relative motion amplitude of the display so that it matched the perceived depth of a hand-held 7 cm block of wood. The required amplitudes were slightly less than would be required to exactly simulate a depth of 7 cm.

Sakuri and Ono (1996) presented observers with displays containing expansion/contraction dot motion, and found that for displays in which the amount of simulated depth was smaller than 4 cm, there was slight over-estimation of the peak-to-trough depth. In a subsequent study, Sakuri and Ono (1997) found depth over-estimation in expansion/contraction displays containing simulated depths of up to 8 cm. In relative motion displays the over-estimation occurred for simulated depths less than or equal to 4 cm.

Durgin et al. (1995) showed observers motion parallax-defined cones, represented by five equally spaced circular contours. The base of the cone was 6 cm. The simulated tip-to-base depth of the cone took a value between 2.4 and 18 cm. The task of the observer was to adjust an icon which represented a side view of the cone, such that the tip-to-base distance of the icon was identical to the perceived tip-to-base depth of the cone. Viewing distance was 72 cm. It was found that for tip-to-base cone depths of less than 6 cm, depth was over-estimated. For cone depths greater than 6 cm, depth was under-estimated. In a second experiment, real, wooden cones were used. The base of the cone was 10 cm. The tip-to-

base depth was 5, 10, 15 or 20 cm. Once again observers adjusted an icon to indicate the perceived depth of the cones. It was found that depth was over-estimated in the two smallest depth conditions, was correctly estimated in the 15 cm condition and was under-estimated in the 20 cm condition.

These results suggest that observers have a tendency to over-estimate the amount of depth when the relative motion amplitudes and thus the simulated depths are relatively small. Other studies, however, have found that depth under-estimation can occur even when the amounts of simulated depths are relatively small. For instance, Rogers and Graham (1979) used a stereoscopic match stimulus to measure the depth perceived by observers in a parallax display. Both parallax and stereoscopic displays subtended 12.5° x 10° and were viewed from a distance of 57 cm. Simulated peak-to-trough depths were 0.66, 1.43 or 3.02 cm. The magnitude of depth was correctly perceived in the stimulus of lowest amplitude, and there was an under-estimation at the two greater amplitudes.

Caudek and Proffitt (1993) also found that observers under-estimated the depth of objects with simulated depths between 20 and 80 mm. The method of depth matching in this study has been described above, and was essentially similar to that used by Durgin et al. (1995). The basis of the respective under-estimation and over-estimation of depth in the two studies is unclear, although a potential factor is the relative widths of the stimuli of each study. Those of the earlier study were smaller than the single width employed by the later study. Since it has been demonstrated that increased amounts of depth are perceived with increasing stimulus width, this may partly underlie the observed differences in results.

Braunstein and Tittle (1988) also found that perceived depths were smaller than simulated, using a dihedral angle stimulus similar to that of Caudek and Proffitt (1993).

It is impossible to pinpoint the reasons for the under-estimations and over-estimations of the amounts of depth perceived in the various studies, but it is notable that depth under-estimation generally occurs for stimuli of relatively small size or high corrugation frequency. For

instance, the corrugation frequency of the surface used by Rogers and Graham (1979) was 0.3 cpd, and the sizes of Caudek and Proffitt's (1993) and Braunstein and Tittle's (1988) stimuli were 2° - 8° and 12° respectively. The corrugation frequencies shown to result in substantial depth over-estimation in Experiment 11 were 0.05 and 0.1 cpd, and were of size 20°.

A further factor influencing the amount of perceived depth is the head movement of the observer: a greater amount is generally perceived in conditions of head movement than in conditions in which the observer remains stationary (Ono and Steinbach, 1990). Most (although not all) of the studies in which over-estimation was observed involved head-linked parallax, and most (although again, not all) of the studies in which under-estimation was observed involved stationary observers.

To summarise and conclude, the current experiment demonstrated that increases in corrugation frequency lead to decreases in perceived depth and to concomitant increases in perceived rotation. That perceived depth is affected by corrugation frequency is consistent with previous work indicating that depth is scaled with the width of a single cycle of a triangle wave (Caudek and Proffitt, 1993) and with work indicating that perceived stereoscopic depth magnitude decreases with increasing disparity gradient (Bülthoff et al., 1991), although the parallels which may be drawn between the latter study and the current one are limited. The observed depth/rotation trade-off is qualitatively consistent with the geometric model, although the combined magnitudes of perceived depth and rotation generally far exceeded those predicted by the model. The following experiment investigates the relationship existing between depth and rotation when stimulus size is manipulated.

5.3 Experiment 12: Matched Rotation and Matched Depth as a function of Stimulus Size

5.3.1 Stimuli

Stimuli were of size 2.5, 5, 10 or 20°. The corrugation frequency was 0.4 cpd. The relative motion amplitudes were identical to those of the previous experiment.

5.3.2 Procedure

This was identical to that of the previous experiment.

5.3.3 Results

Figs. 5.19 - 5.22 show Matched Rotation and Matched Depth as functions of relative motion, for the four stimulus sizes above. As stimulus size increased, Rotation Matches decreased and Depth Matches increased.

Figs. 5.23 and 5.24 show the data fit to the model. The lower legend shows the relative motion amplitude (arc min equivalent disparity) which produces each theoretical depth/rotation curve. The upper legend shows the surface corrugation frequency. Symbols of identical colour arose from surfaces containing identical relative motion amplitudes, the values of which are indicated in the former legend. The data should theoretically lie on the curve of the same colour. It is evident that, for all surfaces except those containing the lowest relative motion amplitude, the combined amounts of depth and rotation are greater than those predicted by the model. The data generally lie closer to the ordinate and are displaced upwards, relative to those of the previous experiment. This indicates that greater amounts of rotation and lower amounts of depth were generally perceived in the current experiment.

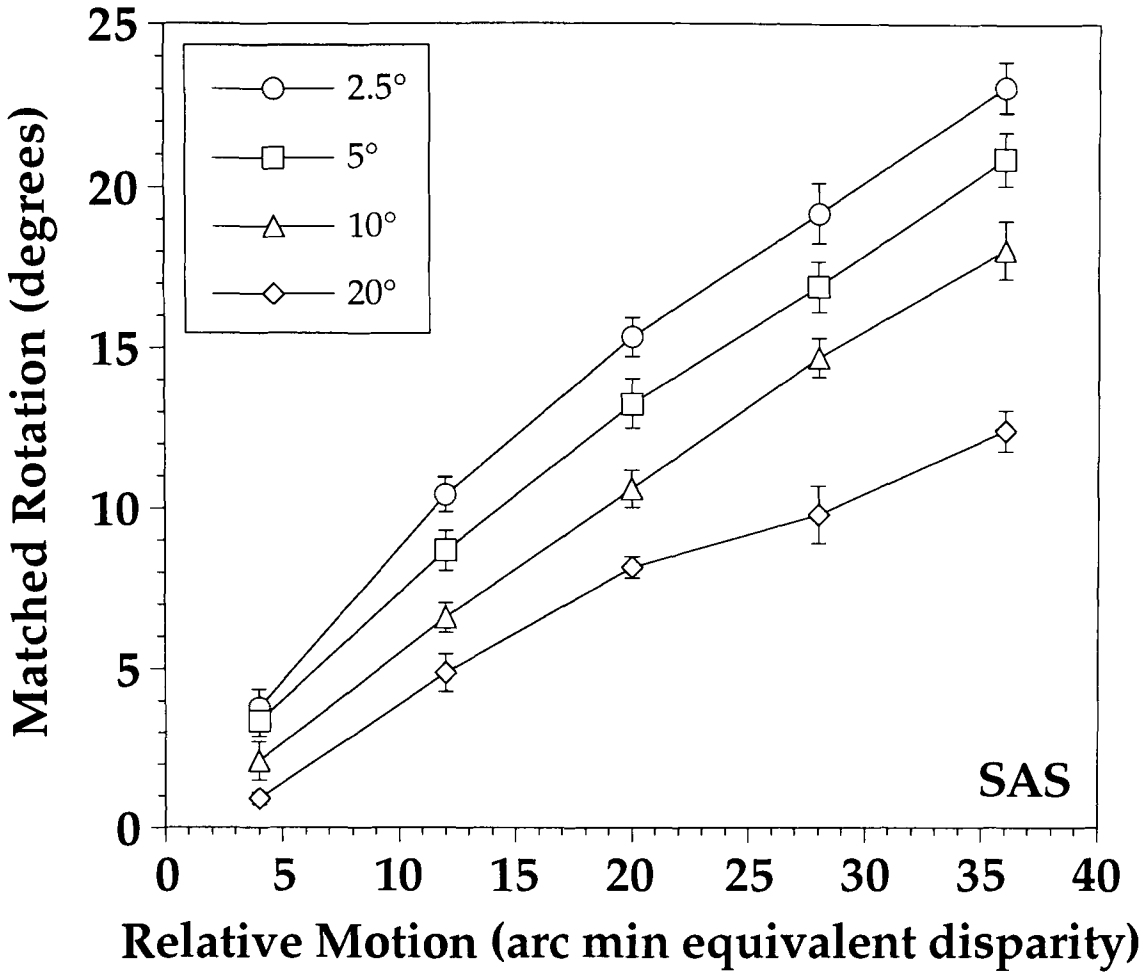


Fig. 5.19: Matched Rotation as a function of Relative Motion for four stimulus sizes, for Subject SAS.

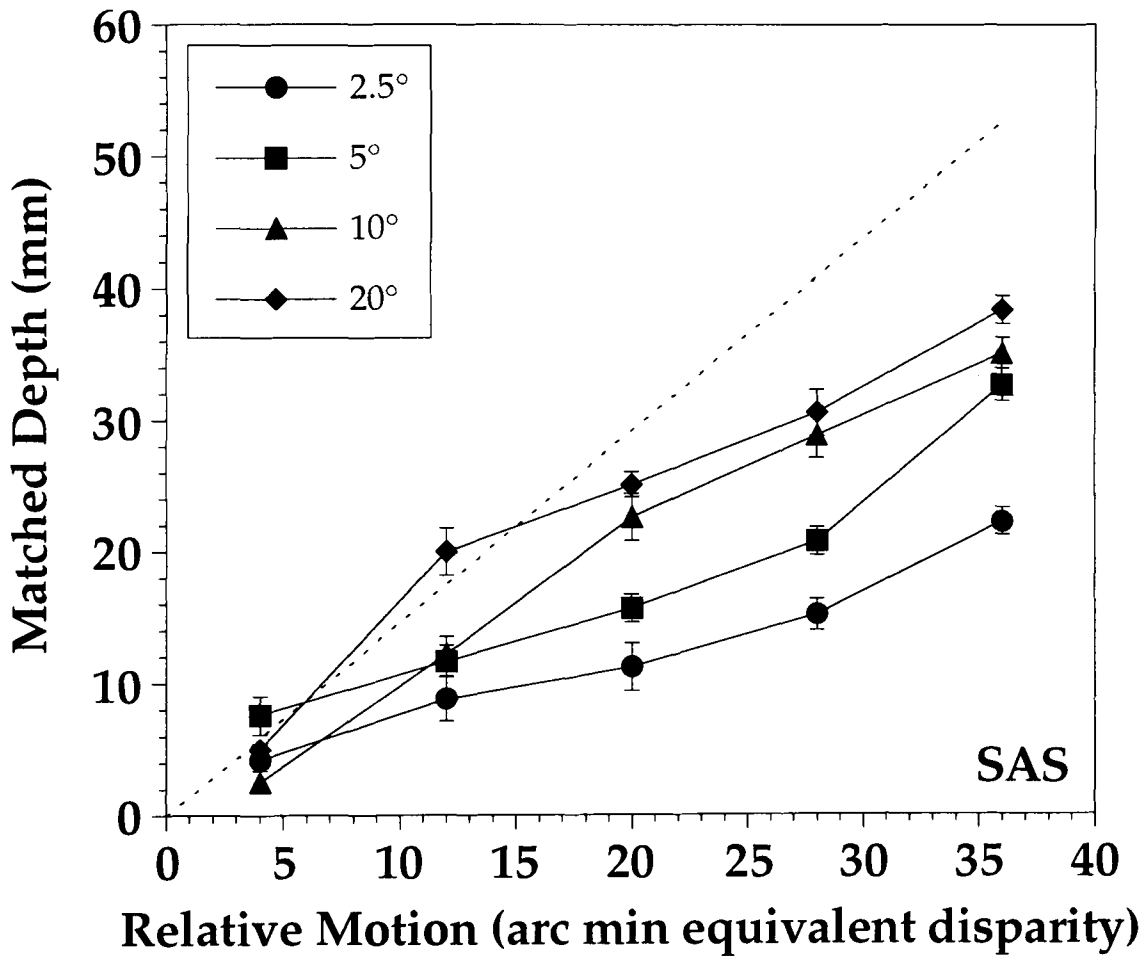


Fig. 5.20: Matched Depth as a function of Relative Motion for four stimulus sizes, for Subject SAS. The dashed line represents simulated depth.

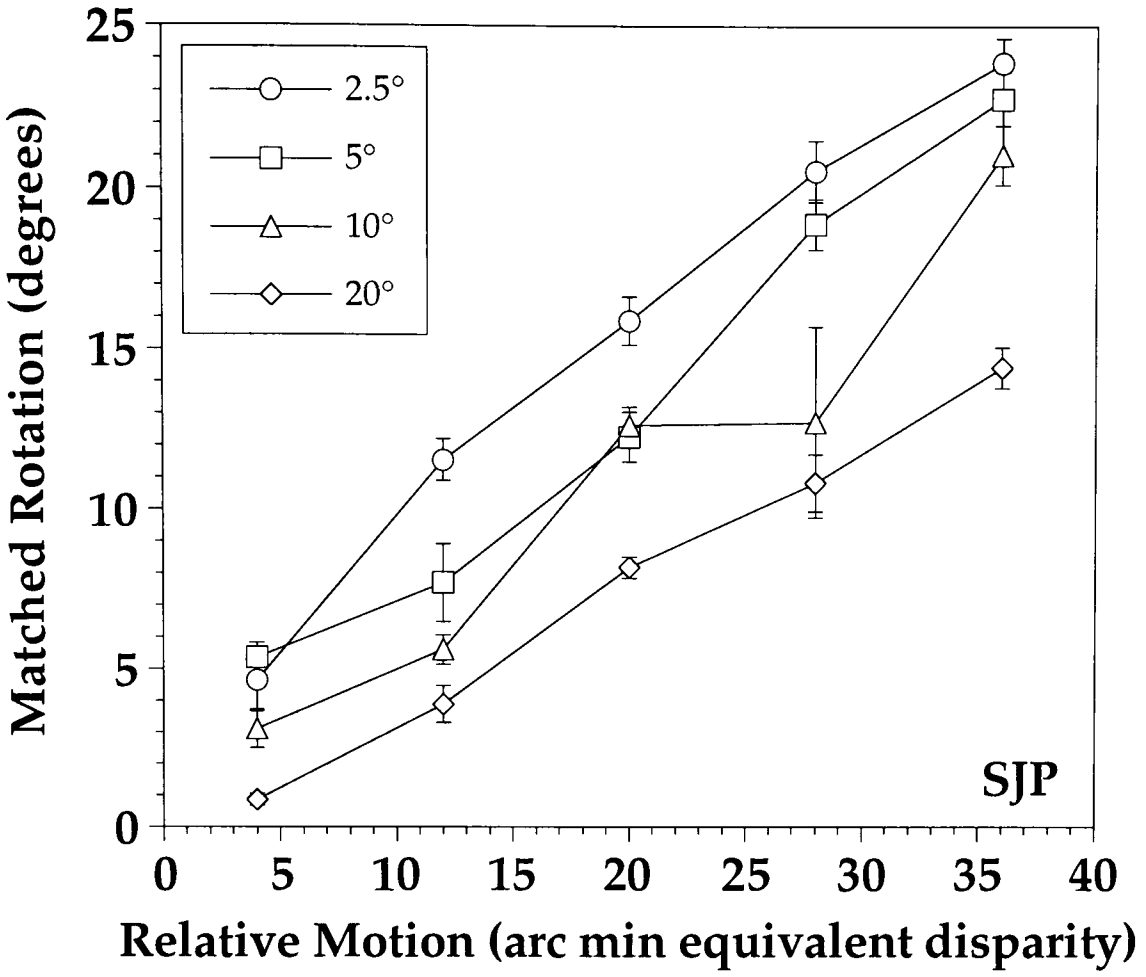


Fig. 5.21: Matched Rotation as a function of Relative Motion for four stimulus sizes, for Subject SJP.

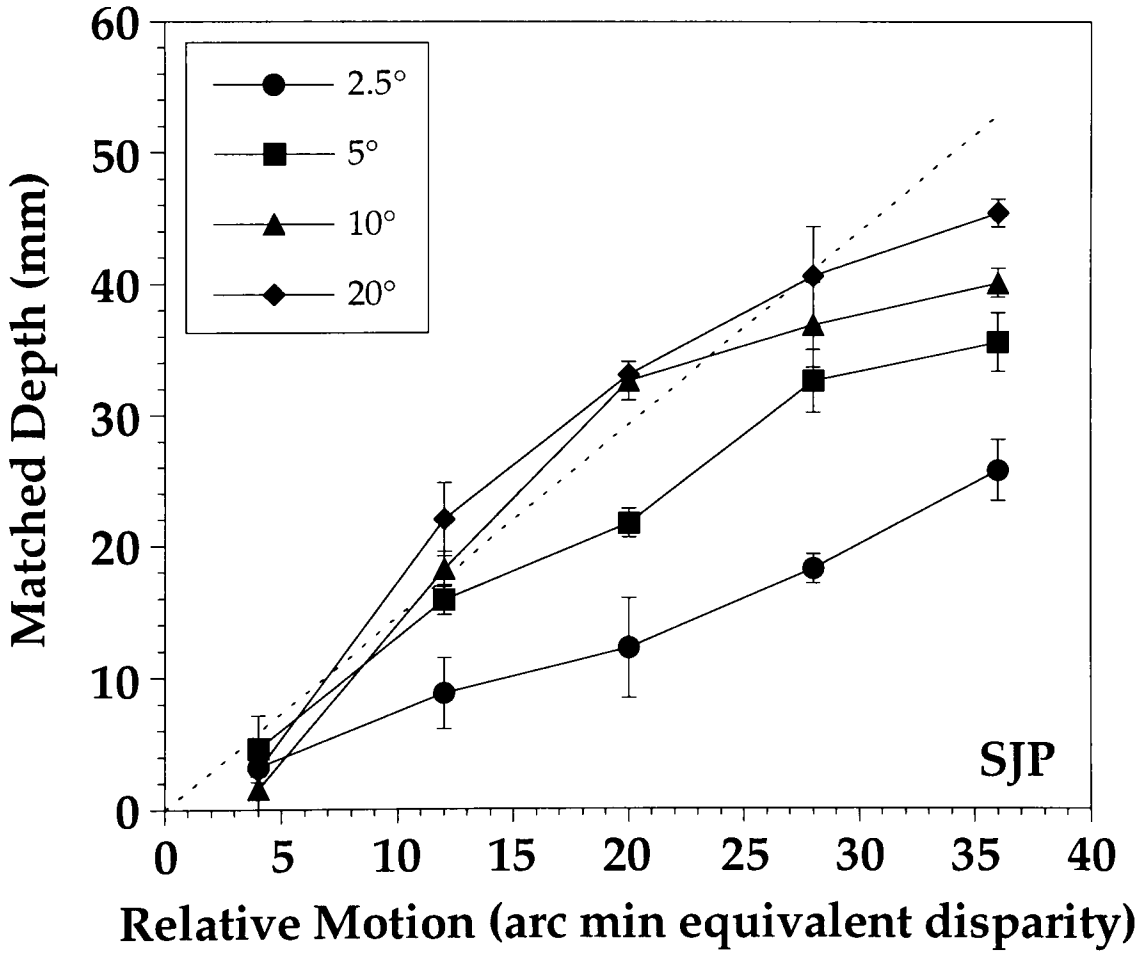
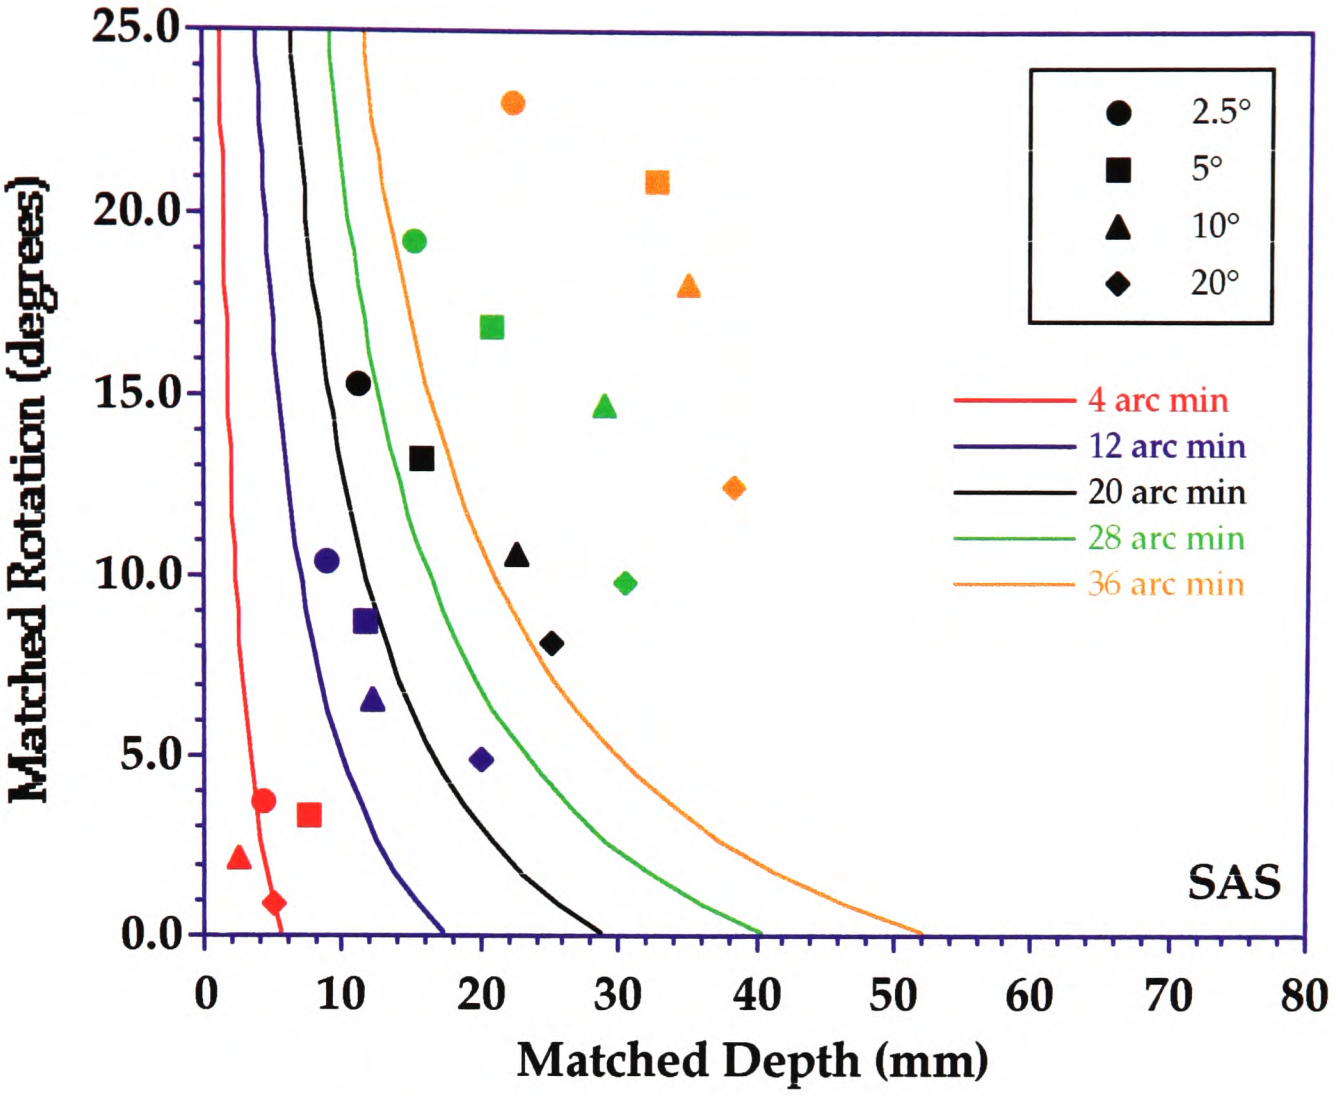
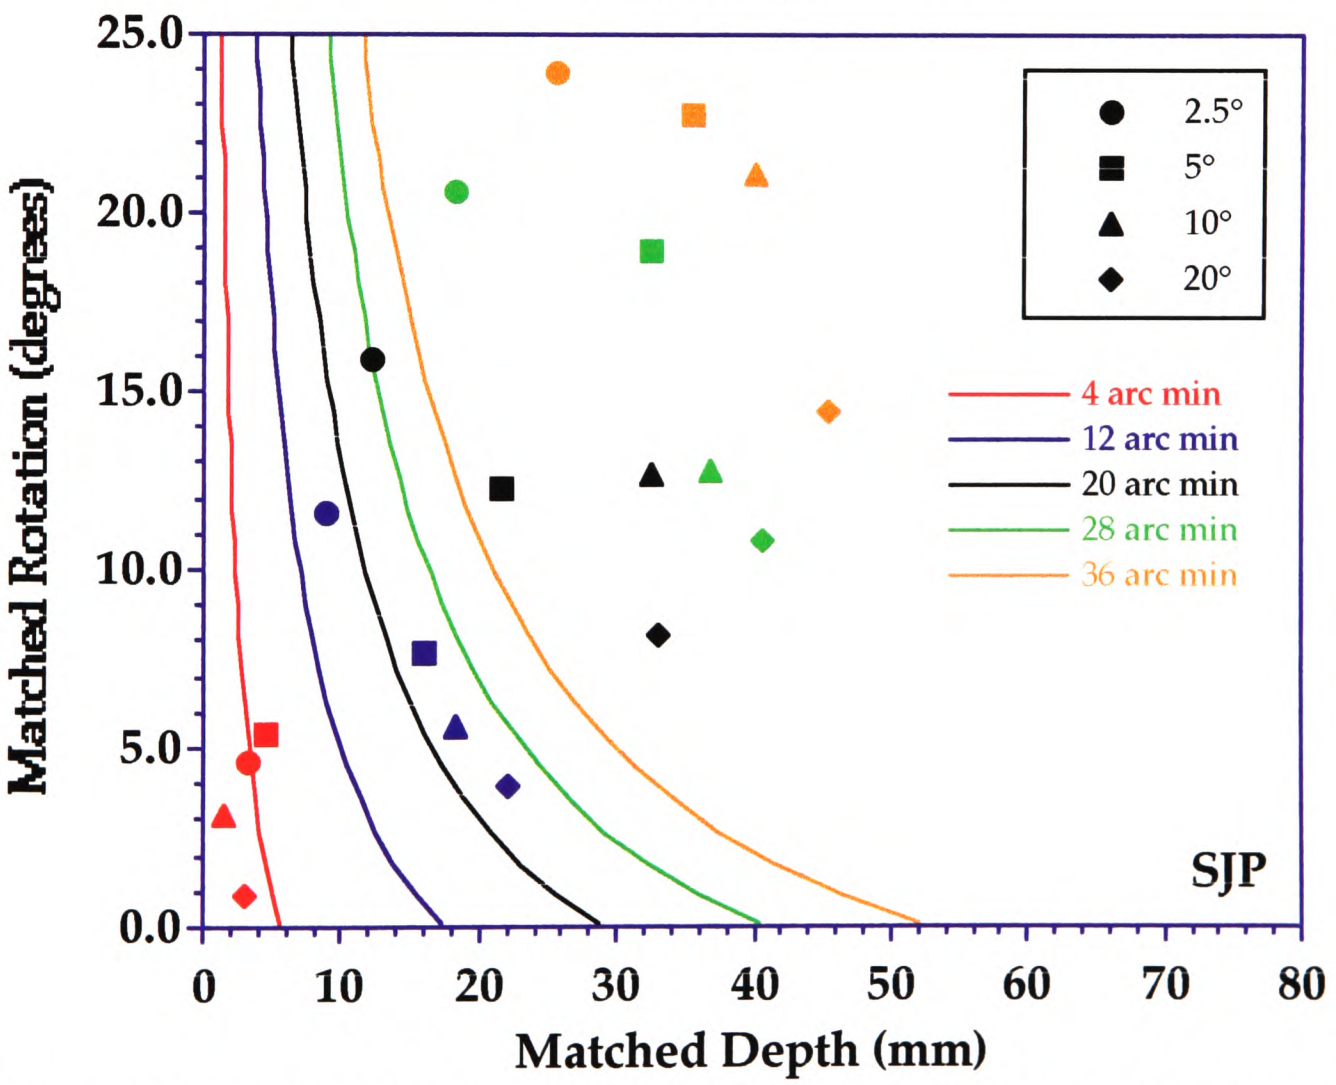


Fig. 5.22: Matched Depth as a function of Relative Motion for four stimulus sizes, for Subject SJP. The dashed line represents simulated depth.





Figs. 5.23 & 5.24: Matched Rotation versus Matched Depth for four stimulus sizes, for Subjects SAS and SJP. The curves show the rotation/depth combinations predicted by the geometric model for the five different relative motion amplitudes. Data should theoretically fall on the line of the same colour.

Figs. 5.25 and 5.26 show the relative motion amplitudes which are theoretically required to produce the observed combinations of depth and rotation magnitudes. The theoretical amplitudes generally lie between 1.5 and 2.5 times the physical amplitudes of the surfaces. The extent of over-estimation is slightly lower in this experiment than in Experiment 11. Differences between the extents of over-estimation are small at the lower relative motion amplitudes, but become more evident as the amplitude increases. At higher amplitudes the over-estimation is generally lowest for the 2.5° surface and highest for either the 5° or 10° surface.

5.3.4 Discussion

The finding that perceived rotation decreases with increasing stimulus size has previously been noted and discussed in Section 4.6. The finding that perceived depth increases with increasing display size is consistent with a trade-off between the respective magnitudes of perceived rotation and perceived depth, and is a result qualitatively similar to that of Experiment 11 above. Figs. 5.23 and 5.24 show that the combinations of perceived depth and rotation are larger than predicted by the geometric model. Once again it would seem that, under these experimental conditions, observers' percepts of depth and rotation magnitudes are *not* quantitatively constrained by the model. Such a conclusion must not be drawn too hastily however, since the point raised in the discussion of the preceding experiment — namely that concerned with the potential over representation of the perceived amounts of depth and rotation — has yet to be evaluated (see Section 5.4).

Figs. 5.25 and 5.26 quantify the relationship between the physical amplitude of each surface and the amplitude of relative motion theoretically required to produce the observed magnitudes of perceived depth and rotation of that surface. As in the previous experiment, the graphs illustrate that the theoretical amplitudes are significantly higher than the physical amplitudes of the surfaces. The discrepancy between the two amplitudes is generally slightly smaller in the current experiment. This is likely to be because the surfaces of the present experiment were generally perceived to contain less depth than those of the previous experiment. The theoretical increases in relative motion amplitude required to support the larger extents of perceived rotation observed in the

current experiment are *less* than the theoretical decreases in relative motion amplitude which support the decreased extent of perceived depth. This is suggested by the shape of the model's depth/rotation curves. A given relative motion amplitude supports an increase in perceived rotation from a moderate level to sizeably greater level without appreciable decreases in the extent of perceived depth. Large increases in perceived depth, however, are necessarily accompanied by substantial decreases in perceived rotation. Thus in order to perceive sizeable increases in depth accompanied by only moderate decreases in rotation, a greater relative motion amplitude is required. In sum, moderately-sized changes in the magnitude of perceived rotation do not have such a great effect upon the theoretical relative motion amplitudes as do moderately-sized changes in the magnitude of perceived depth.

The finding that increased amounts of depth are perceived with increasing display size is inconsistent with the results of Ono et al. (1986). In Experiment 1 of that study, observers were presented with displays containing 28.2 arc min equivalent disparity. A greater amount of depth was perceived in a display which subtended 20° x 21° than in a display which subtended 40° x 42°. Furthermore, the corrugation frequencies of the displays were 0.08 and 0.04 cpd respectively — an additional factor which, on the basis of the results of Experiment 12 of this thesis, would predict a greater amount of perceived depth in the larger stimulus.

In Experiment 2 of Ono et al. (1986), subjects adjusted the amount of motion parallax in the display until the perceived depth was identical to that of a hand-held, 7 cm block of wood. It was found that a slightly greater amount of relative motion was required for the large display than for the small display. This suggests that a given amount of relative motion produced a greater amount of perceived depth in the small than in the larger display.

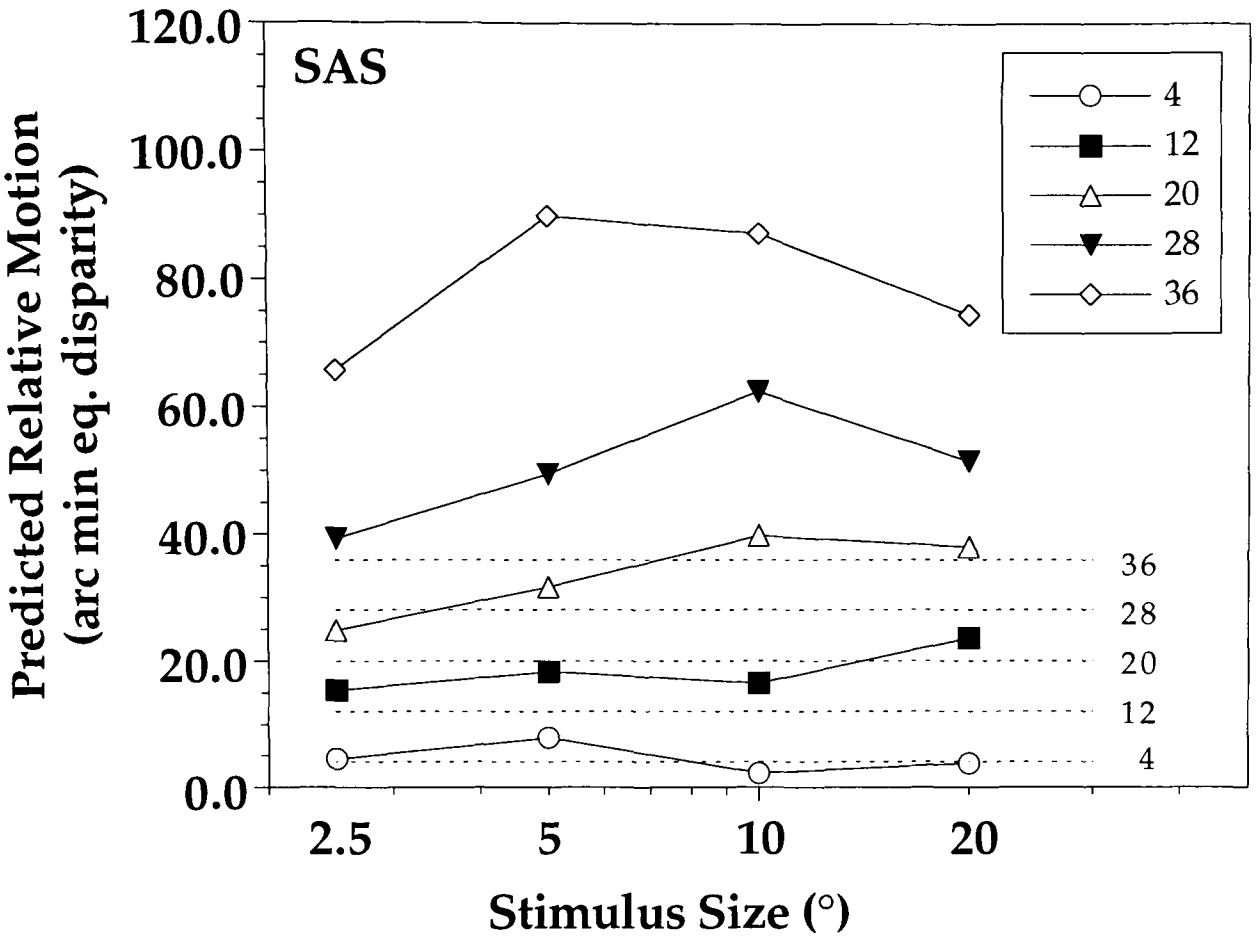


Fig 5.25: The relative motion values which would theoretically be required to produce the perceived amounts of rotation and depth in each surface, for Subject SAS. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

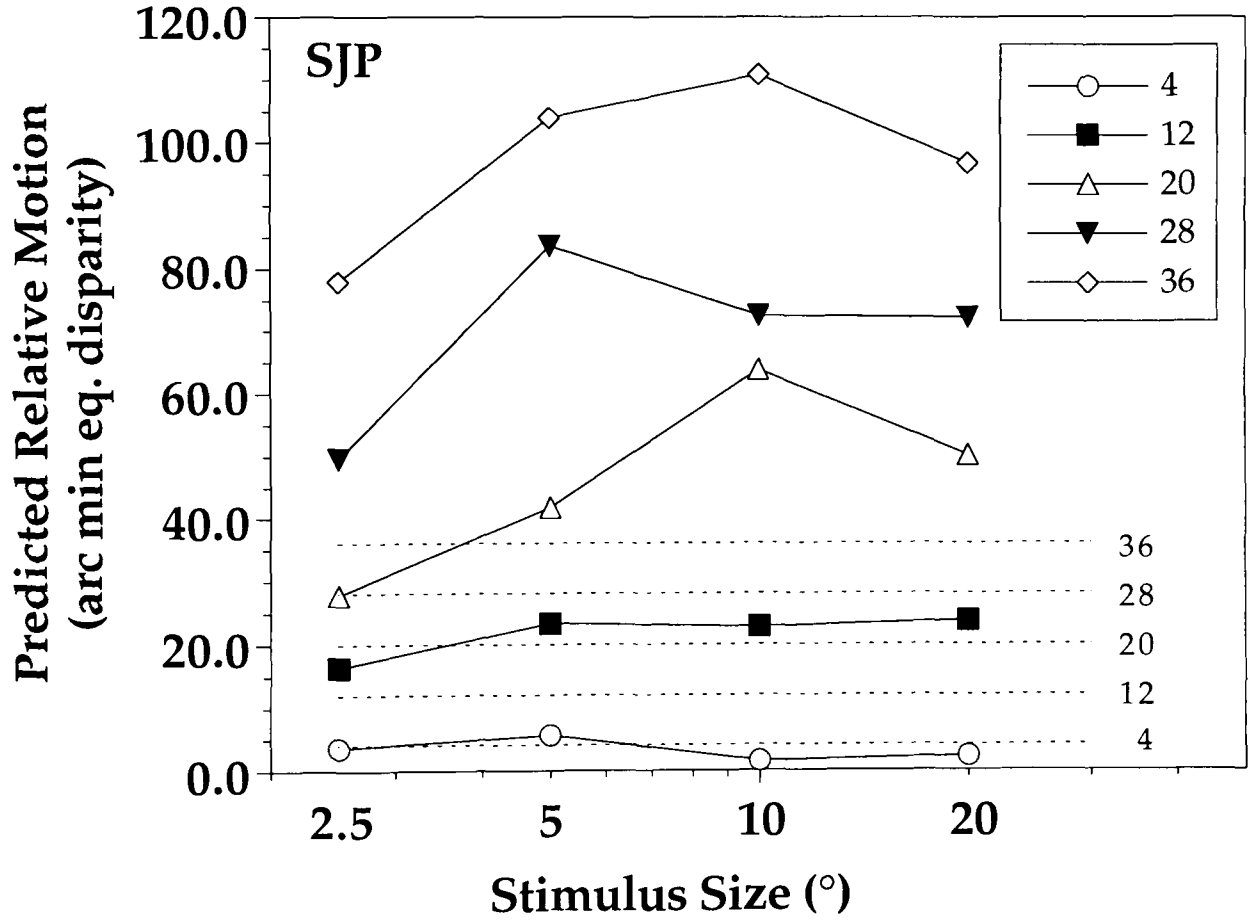


Fig 5.26: The relative motion values which would theoretically be required to produce the perceived amounts of rotation and depth in each surface, for Subject SJP. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

The basis of the differences between these results and those of Experiments 11 and the current experiment is unclear. It must be noted however that the effects of display size in the study of Ono et al. (1986) accounted for less than 2% of the variance in each experiment. The main purpose of that study was to investigate the effect of viewing distance on perceived depth, and this effect accounted for the vast majority of the variance. The authors themselves did not have an explanation for the effect of display size on perceived depth.

The depth matching results of the present experiment indicate that observers generally tended to under-estimate the amount of simulated depth. This is in contrast to Experiment 11, where a large degree of over-estimation was observed. It is also evident, however, that the extents of perceived rotation were considerably greater in the current experiment. This is likely to account for the generally lower perceived depths. The single surface which appeared in both experiments, i.e. the 0.4 cpd, 20° surface, produced depth and rotation matches which were approximately similar in both experiments.

The current results and those of Experiment 11 demonstrate that both an increase in corrugation frequency and a decrease in stimulus size lead to an increase in perceived rotation and simultaneously to a decrease in perceived depth. This might lead to the conclusion that altering the amount of perceived rotation (for example) of a given stimulus *necessarily* leads to an alteration in the amount of perceived depth, and that the two percepts are causally related. This notion stems from the geometry of the situation: the relationship between these two attributes as specified by the model is indeed causal.

The above conclusion would predict that sizeable changes in the perceived magnitude of rotation would go hand in hand with changes in the magnitude of perceived depth, regardless of the factor(s) causing the changes in the perceived rotation magnitude. This conclusion is not necessarily true, however. It may be the case that the manipulation of corrugation frequency or stimulus size causes changes in the depth and rotation percepts which are specific to those particular manipulations. This possibility is investigated in Experiments 16 and 17, which use a vertical perspective cue to manipulate perceived rotation.

5.4 General Discussion

Both Experiment 11 and Experiment 12 produced data which conformed qualitatively to the depth-rotation trade-off model: manipulation of either the size or the corrugation frequency of a parallax surface resulted in increases in the perceived depth magnitude which were accompanied by decreases in the perceived rotation magnitude, or vice versa. This result suggests that the visual system applies a simple geometric constraint when presented with a given extent of parallax motion. It complements the findings of Rogers and Collett (1989) and Williams (1993). Both of these studies employed composite stereo-parallax surfaces. Both found that perceived depth magnitudes were largely influenced by the disparity information; when the motion parallax amplitude was discrepant with the disparity amplitude, a portion of the former was interpreted as rotation in order that it 'fitted in' with the depth magnitude suggested by the former cue. Both studies concluded that the evoked percept was one which minimised the discrepancies between the various cues available to the visual system concerning the depth and rotational characteristics of the surface.

As noted, these studies found that the perceived depth magnitude of a surface was largely influenced by the disparity information but that its combined depth and rotational qualities were governed, qualitatively at least, by simple geometric considerations. The results of Experiments 11 and 12 demonstrate that simple geometric principles are also used by the visual system in the absence of disparity information — when the system is presented with motion parallax alone.

Over-estimation of depth and rotation

In all cases there was an over-estimation of rotation for the amount of depth perceived, or vice versa. The over-estimation of depth and rotation for a given amount of relative motion has also been observed in other studies.

Williams (1993, Experiment 5) presented observers with 0.28 cpd surfaces specified by both motion parallax and binocular disparity. The amplitudes of each cue took values of 0, 40 or 80 arc min (equivalent)

disparity. These were combined in all possible ways, giving a total of eight surfaces (the surface in which both cues had zero amplitude was not used). Subjects made a rotation match and a depth match. The former match was carried out using adjustable angled paddles, placed 6.5 cm apart. The subject could adjust the slant of these, and the difference between the slants of the two paddles was taken as the extent of perceived rotation. The depth matching was carried out using a stereoscopic sinusoidal matching surface.

The data lay fairly close to the predictions of the geometric model, although there was over-estimation of the combined amounts of depth and rotation for almost every surface, for all subjects. The perceived depth of surfaces containing a non-zero disparity cue was generally close to the depth specified by the stereo information (consistent with the results of Rogers and Collett (1989)), with the over-estimation being due to a greater than predicted perceived rotation magnitude for the amount of depth perceived.

The matched depth magnitude of surfaces containing zero binocular disparity lay between 25 and 35 arc min when the parallax amplitude was 20 arc min equivalent disparity, and was about 35 arc min when the parallax amplitude was 40 arc min equivalent disparity. The magnitudes of perceived rotation were significantly greater than predicted for these perceived depths. It is obvious that the perceived depths were much less heavily constrained by the stereo information when the disparity amplitude was zero. The over-estimation of depth is similar to that observed in Experiments 11 and 12, although the over-estimation was slightly less in the former experiment, perhaps because of the presence of the zero binocular disparity cue.

In subsequent experiments Williams (1993) manipulated the magnitude of perceived rotation using a dynamic perspective cue. In a 'typical example case' (p. 234) a surface containing a relative motion amplitude of 20 arc min equivalent disparity and a binocular disparity of 3.81 arc min, with the perspective cue indicating a rotation angle of 50°, was perceived by Subjects JSW and BL to undergo rotations of 45° and 58° respectively. The matched depths were 8 and 7.5 arc min disparity. The relative motion amplitudes which would theoretically be required to

produce these depth/rotation combinations are 55 and 60 arc min equivalent disparity. These data illustrate a significant over-estimation of the combined depth/rotation values relative to those predicted by the model.

Studies using surfaces defined by parallax alone also provide evidence that the combined extents of depth and rotation are over-estimated in certain circumstances. Ono and Steinbach (1990) asked observers to report the amount of perceived depth and perceived motion in parallax stimuli of various velocity gradients. Two extents of relative motion, simulating depths of 0.4 and 4.0 cm, were used. The trade-off between depth and rotation was found as predicted, but also evident from their data is the fact that relative motion was interpreted both as depth *and* as motion. For example, a sine waveform containing the larger extent of relative motion was perceived to contain almost exactly 4 cm of depth whilst also containing dots which appeared to translate through 1 cm. Similarly, a square waveform appeared to contain about 2 cm of depth whilst containing dots which appeared to translate through 4.5 cm. Ono and Steinbach (1990) do not compare the absolute values of their depth and motion data with the depth/rotation combinations allowed by geometric models since their study was concerned with relative, rather than absolute, magnitudes of perceived depth and motion. However, it is evident that significantly greater extents of relative motion than those used in the study would be required to produce the observed magnitudes of depth and rotation were it the case that relative motion could be perceived *either* as depth *or* as rotation.

Ono, Rivest and Ono (1986) asked observers to estimate the perceived depth in parallax stimuli containing relative motion which simulated either 2.3, 4.9 or 11.2 cm at the 40 cm viewing distance, and additionally to indicate whether the surface appeared to 'rock' about a vertical axis. It was found that, at least for the two lowest values of relative motion, the amount of depth perceived by some subjects was greater than the amount simulated, and furthermore, that rocking motion was also perceived.

Sakuri and Ono (1996) used displays in which depths of 1 - 16 cm were simulated by an expansion/contraction motion. Observers indicated the magnitude of perceived depth by adjusting the distance between two

pointers, and reported verbally whether a 'rocking' motion was perceived. Surfaces with a simulated depth of 4 cm had a perceived depth of 4 cm and were also perceived to rock, indicating that in this condition the relative motion was indeed perceived as both depth and motion.

The instances cited above indicate that under some circumstances the amounts of perceived depth and rotation are greater than those predicted by any model which allows relative motion to be interpreted *either* as depth *or* as rotation. In Section 5.2.4, three possibilities for this apparent over-estimation were suggested. These are considered in turn, below.

Over-estimation of the extent of 2-D relative motion

Firstly, it might be the case that there is indeed a trade-off between the magnitudes of depth and rotation which is quantitatively constrained by the geometric model, but that the visual system over-estimates the extent of 2-D relative motion and thus produces correspondingly large estimates of the stimulus' 3-D attributes. Evidence suggests that 3-D motion perception is constrained by 2-D motion perception (Bradshaw et al., 1991; Eagle and Blake, 1995; Werkhoven and van Veen, 1995; Hogervorst and Eagle, 1998; Ichikawa and Ono, 1998) (although see van Damme and van de Grind, 1996). If the latter were over-estimated then one would expect the former to be over-estimated also. Over-estimation of 2-D motion has been reported by Sato (1989) and by Ono and Steinbach (1990). Sato found that a small displacement in a random dot kinematogram produced an exaggerated extent of perceived motion. Ono and Steinbach (1990) found that stationary observers viewing a square wave relative motion display perceived 65 mm of relative dot motion (and no depth) when the true relative motion amplitude was 15 mm.

An over-estimation of the relative motion could arise, for example, through an over-estimation of the viewing distance. In a stereoscopic shape judgement task carried out by Johnston (1991) (see Section 6.5.5), depth was over-estimated at a viewing distance of 53.5 cm. Johnston suggested that her results were consistent with a default assumed viewing distance of 80 cm. Previous studies have also concluded that viewing distance is often not veridically perceived (e.g. Foley, 1980). Durgin et al. (1995) found near-perfect stereoscopic depth constancy however, and

suggested that the reason for previous studies' failure was the low illumination levels generally used. Evidence suggests that distance perception is faulty in these conditions or in the absence of a surrounding context (e.g. Owens and Leibowitz, 1980). Gogel and Tietz (1973) have shown that in dim light observers tend to over-estimate the distance of near targets and under-estimate the distance of far targets, a bias termed the 'specific distance tendency'. Objects in dim lighting were generally perceived to lie at a distance of about 2 m (Gogel, 1982).

In a well-lit, fully structured environment Glennerster et al. (1996) found that stereoscopic depth constancy was dependent upon the task: performance was poorer when subjects were required to make the judgements described above — namely the adjustment of a hemi-cylinder to apparent circularity and the adjustment of a triangle wave to 90° (note, however, that performance was approximately 70%, a value considerably higher than that found by Johnston (less than 30%)) — but was near-perfect when the task merely involved depth matching. The authors suggested that the former task might involve the construction of a metric representation of the surfaces, for which an accurate estimate of the viewing distance was required, whereas a simpler 'direct' strategy which did *not* require information about the viewing distance might be employed for the latter. This would not, however, explain the near-perfect performance of the subjects of Durgin et al. (1995) on a task which also required knowledge of the viewing distance. These authors suggest that perception of distance approaches veridicality at close viewing distances in well lit environs (p. 692).

In the case of parallax, inaccurate estimates of viewing distance *cannot* be used to account for the relatively poor performance, at least in the study of Durgin et al. (1995). Stereoscopic tasks requiring such an estimate were performed extremely well whereas the same tasks performed with parallax stimuli were not. Of course, this does not rule out the possibility that viewing distance was not over-estimated in Experiments 11 and 12, but it lessens the likelihood that perceived distance was a factor, especially since the surrounding environment was visible and well structured.

An over-estimation of the amount of relative motion could also arise through an under-estimation of the extent of head movement. However,

studies by Gogel (1982) and Gogel and Tietz (1979) suggest that the perceived and actual extents of head motion are equal.

Interpretation of a given extent of relative motion as both depth and rotation

The second possibility is that observers perceive a veridical amount of relative motion, but they interpret the motion both as depth *and* as rotation. A trade-off between the percepts occurs, such that increases in the perceived rotation occur concomitantly with decreases in perceived depth and vice versa, but for any given value of relative motion the combined magnitude of these attributes is greater than predicted by the model.

The explanation of the over-estimation may be similar to the idea discussed in Section 4.3.5 — namely that a given amount of relative motion may give rise to qualitatively different percepts depending upon the nature of the task. When for example a depth judgement is made, attention is directed towards that particular surface attribute and it is usually the case that much of the relative motion is interpreted as depth. If a rotation judgement is subsequently made, the amount of relative motion contributing to that percept may be greater than the amount 'unused' by the depth percept simply because the relative motion is indeed ambiguous and attention is currently directed toward the rotation percept. This argument would seem to suggest that the perceived magnitudes of depth and rotation change as attention is shifted between them, a notion which on the face of it seems untenable since one would expect the observer to perceive a surface the depth and rotational characteristics of which remain constant. However, if while attention is focused on a given attribute the magnitude of the other attribute is not consciously quantified, then it is entirely possible that a given amount of relative motion can 'count' both as depth and successively as rotation (or vice versa). It seems unlikely that the extent of over-estimation would be as large as that observed in Experiments 11 and 12, but in order to eliminate this possible explanation of the over-estimation, a task involving simultaneous depth and rotation judgements should be used. Such a procedure was used by Van Veen and Werkhoven (1996): observers simultaneously adjusted the slant and the rotation of a planar match surface in order to make it appear identical to a test surface. Their data (Fig. 2, p. 2202) show that there was

approximately as much under-estimation of the combined attributes as there was over-estimation. Note though, that their stimuli were not the same as those employed in this thesis.

Biases due to the psychophysical procedures employed

The possibility that the apparent over-estimations are due to the rotation and depth matching techniques themselves must also be considered. Digital callipers have previously been regarded as an adequate device to measure the perceived peak-to-trough depth and the perceived peak-spacing in a stereoscopic display (Bradshaw, Glennerster and Rogers, 1996). A number of motion parallax experimenters investigating perceived depth have asked observers to adjust the distance between two rods to match the perceived depth in the display — a technique essentially similar to the adjustment of calliper jaws. Obviously, the use of this apparatus and experimental method in previous studies does not validate their use in the experiments of this thesis, the point is merely being made that this procedure has generally been found to be satisfactory.

There appears to be no evidence which suggests that observers are inherently more likely to over-estimate the amount of perceived depth if they use a technique of this nature. In fact, only two of the eight surfaces in Experiments 11 and 12 contained a larger amount of perceived depth than simulated: it was the *combined* depth and rotation magnitudes which were greater than those predicted by the model.

The rotation matching technique used in this thesis has not been previously documented, and thus neither has its efficacy. However, two factors would suggest that observers should find it relatively easy to make veridical rotation matches using this apparatus. Firstly, the match was carried out with a real, rotating paddle, the slant of which was signalled by a range of visual cues. Secondly, the paddle rotated in synchrony with the test surface as well as lying on the same vertical axis as the surface. Observers seemed to find the task relatively simple and intuitive.

It seems unlikely (although it is, of course, possible) that biases inherent in the procedure would result in over-estimations which were quite so large as some of those observed here. Although such biases may

have played a role, it is possibly the case that all three reasons outlined above — namely an over-estimation of the extent of 2-D relative motion, a propensity to perceive a given amount of relative motion both as depth *and* as rotation, and methodological factors — contributed to the large values obtained.

5.5 Experiment 13: Matched Shearing and Matched Depth in Square Wave surfaces as a function of Corrugation Frequency

5.5.1 Introduction

Experiments 11 and 12 demonstrated that a trade-off relationship exists between the perceived depth and rotation magnitudes of a sinusoidal surface. The present experiment investigates whether the same trade-off occurs between the perceived magnitudes of depth and shearing in square wave surfaces. Additionally, it examines whether any such relationship is in quantitative accordance with the geometric model of Fig. 5.8.

Previous work in this area is limited. Ono and Steinbach (1990) measured the amounts of shearing and depth in a square wave surface, and noted that the perceived magnitudes were in a qualitative trade-off relationship under conditions of head movement and no head movement. A 0.1 cpd surface (size 10° x 13°) of relative motion amplitude 12 or 113 arc min equivalent disparity, was used. The present experiment investigated the effect of corrugation frequency on the putative trade-off. Experiment 2 showed that the Transition Point of square wave surfaces decreases with increasing corrugation frequency. Although the Transition Point relates specifically to the perceived *rotation* of the surface, it might be expected that shearing might also be more readily perceived as frequency increases. If this were the case, then concomitant decreases in the extents of perceived depth would be predicted.

5.5.2 Stimuli

Stimuli were depth modulated according to a square waveform motion gradient. Pilot experiments established that the effect of corrugation

frequency on the perceived extents of shearing and depth was relatively small. Therefore only the highest and lowest frequencies of Experiment 11 — 0.05 and 0.4 cpd — were used. Stimulus size was 20°. The relative motion amplitudes were identical to those of Experiments 11 and 12.

5.5.3 Procedure

On each trial, observers judged the perceived magnitudes of depth and shearing between the nearer and the further surfaces of the square wave. In each case observers adjusted a pair of digital callipers to match the perceived extent of the attribute. In half the total number of trials the depth judgement was made first and the shearing judgement was made second. In the other half, this order was reversed. The experimental design was identical to that of Experiment 11.

5.5.4 Results

Figs. 5.27 and 5.29 show matched depth as a function of relative motion amplitude for the two corrugation frequencies, for Subjects SAS and SJP respectively. It can be seen that a greater extent of depth was perceived in the surface of lower corrugation frequency. As the relative motion amplitude increased beyond 20 arc min equivalent disparity, Subject SJP perceived a decreasing amount of depth. Figs. 5.28 and 5.30 show the magnitude of perceived shearing as a function of relative motion amplitude for the two corrugation frequencies, for Subjects SAS and SJP respectively. At the higher relative motion amplitudes, the extent of perceived shearing was slightly lower in the 0.05 cpd surface.

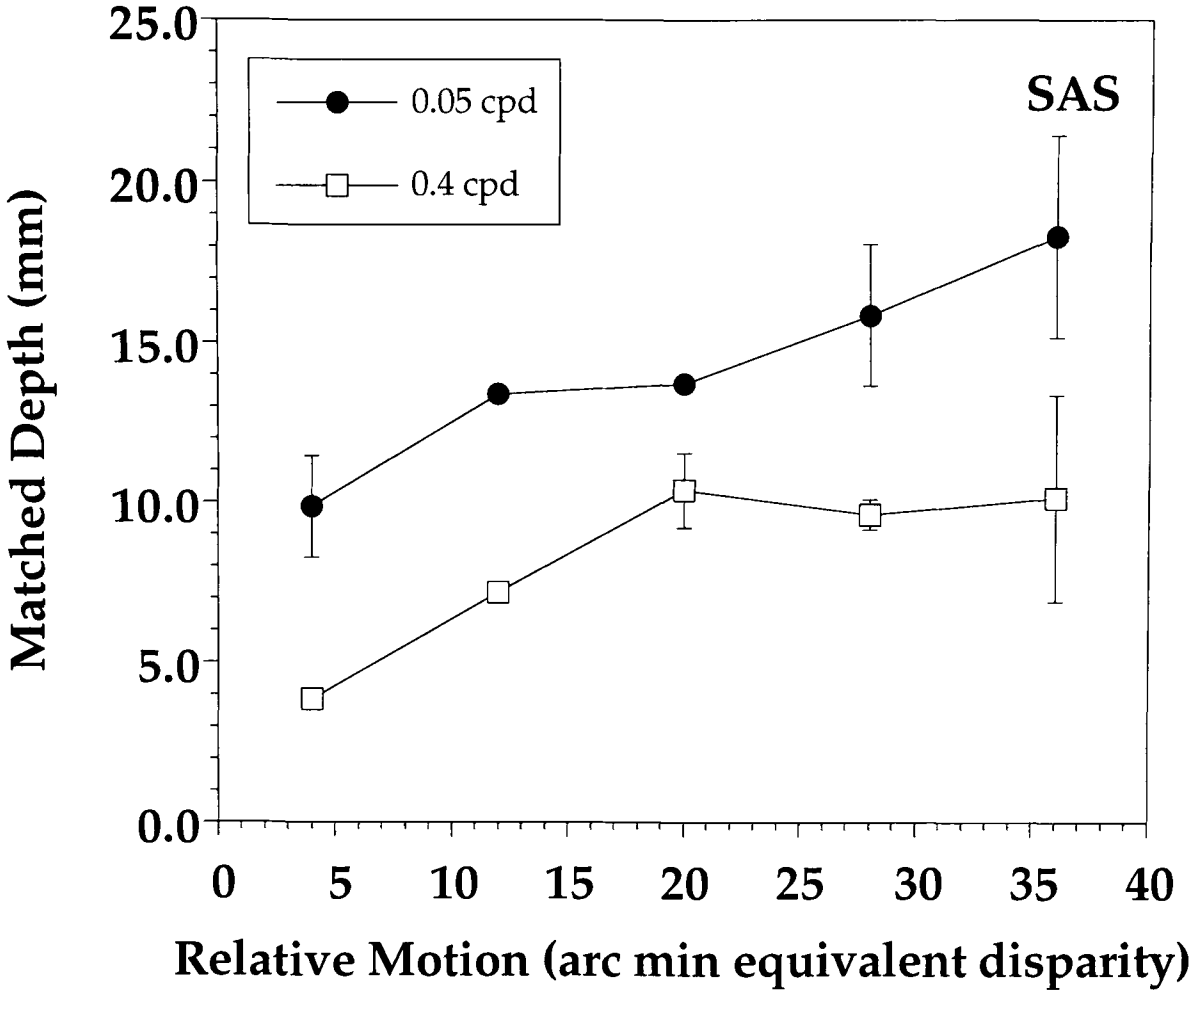


Fig. 5.27: Matched Depth as a function of relative motion amplitude, for Subject SAS.

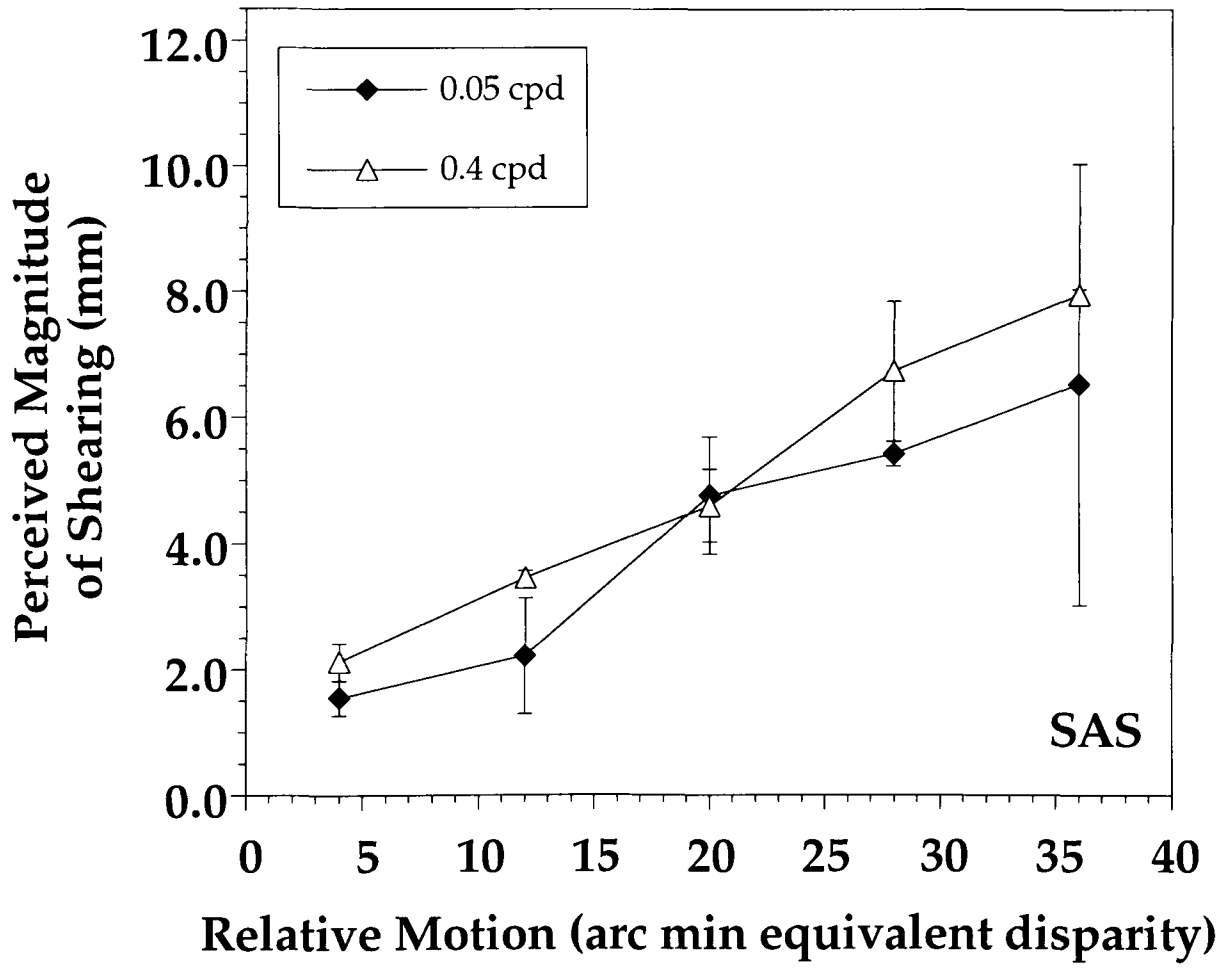


Fig. 5.28: Matched Shearing as a function of relative motion amplitude, for Subject SAS.

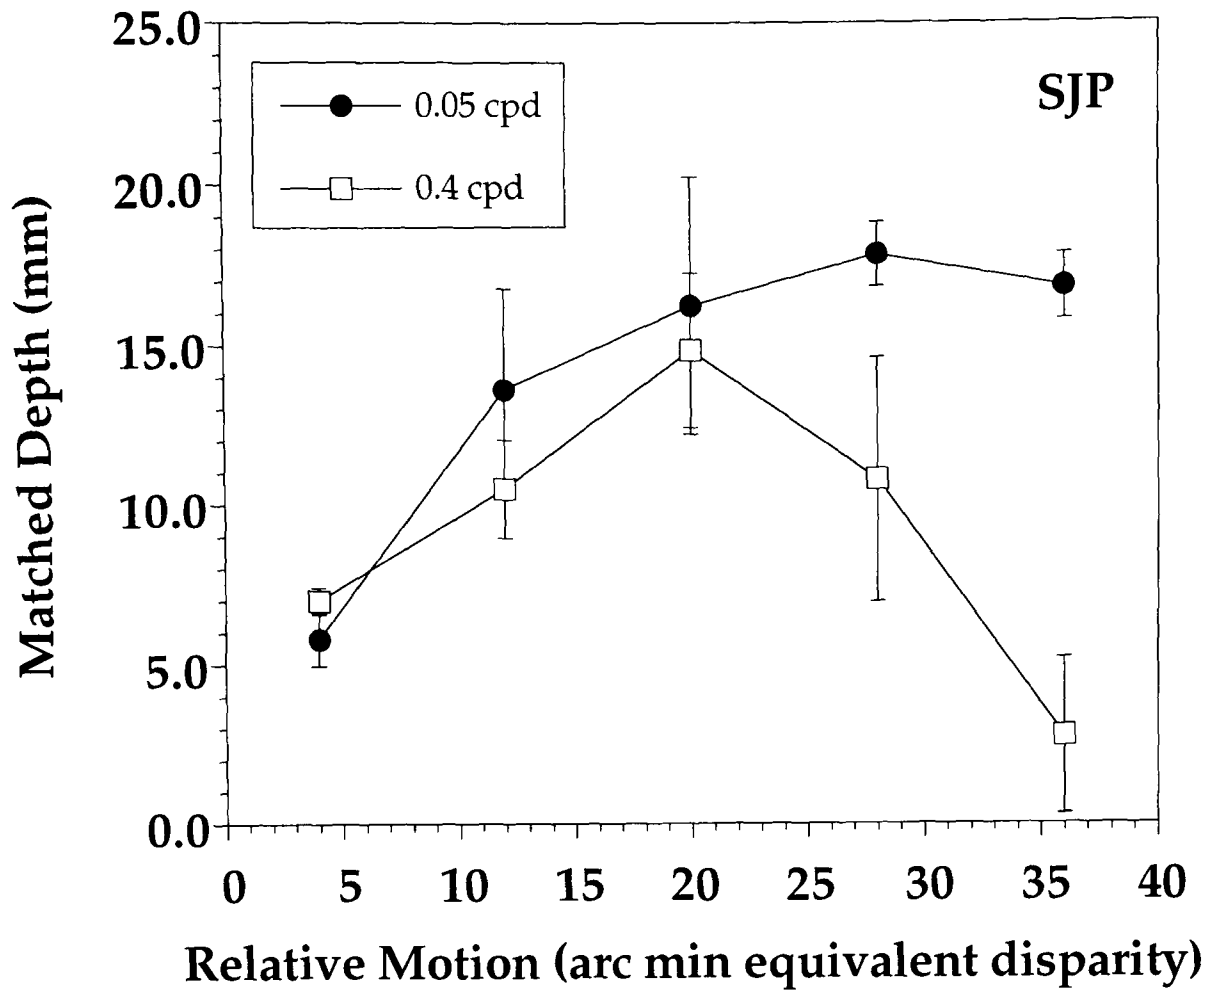


Fig. 5.29: Matched Depth as a function of relative motion amplitude, for Subject SJP.

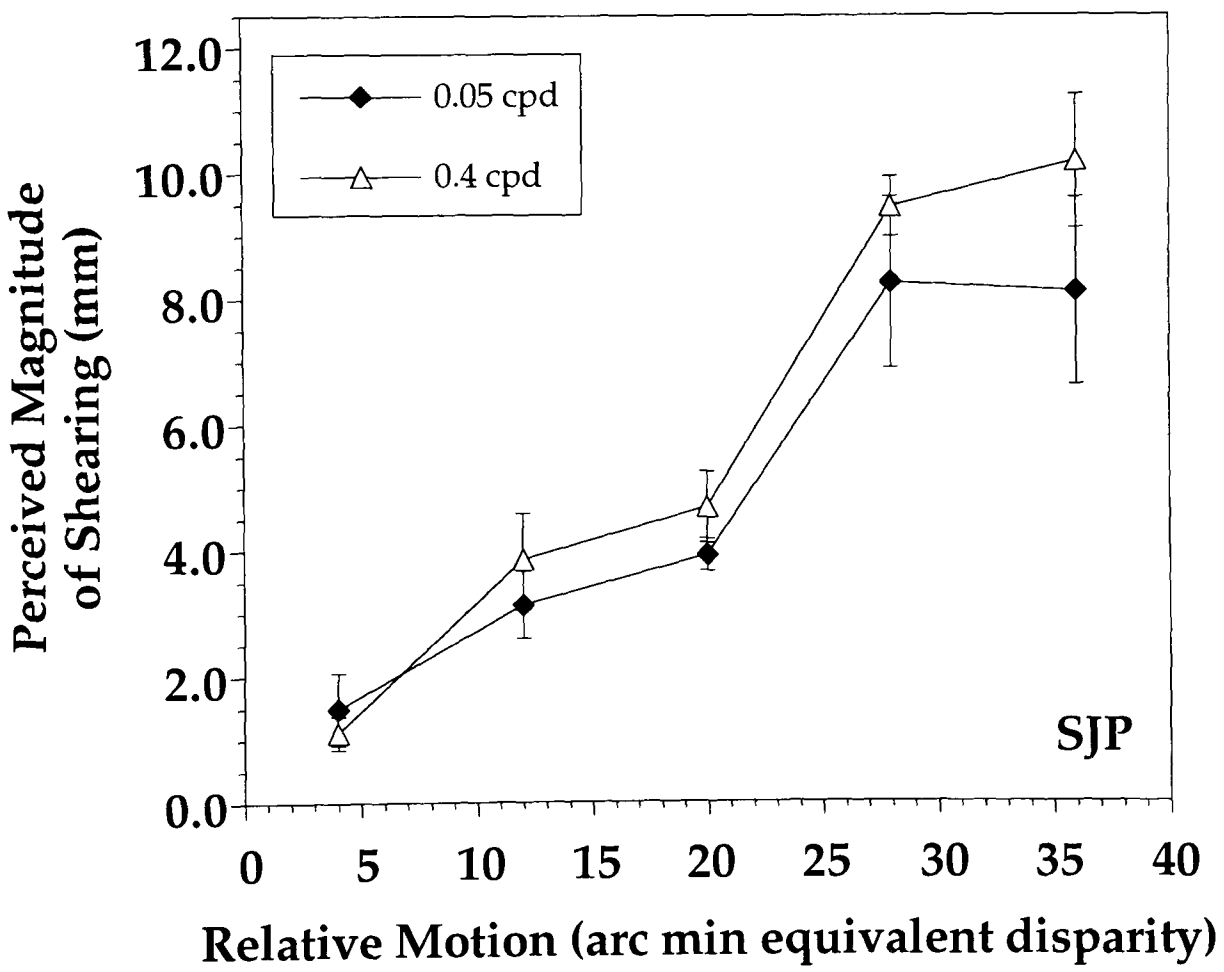


Fig. 5.30: Matched Shearing as a function of relative motion amplitude, for Subject SJP.

Figs. 5.31 and 5.32 replot the data to show its relationship with the

depth/shearing combinations predicted by the model for each amplitude of relative motion. The coloured legend represents the relative motion amplitudes (arc min equivalent disparity) which produce each theoretical depth/shearing curve. The right-hand legend shows the surface corrugation frequency. Symbols of identical colour arose from surfaces containing identical relative motion amplitudes. The data should theoretically lie on the curve of the same colour. It is evident that, for all surfaces tested, there was an over-estimation of the combined amounts of depth and shearing relative to those predicted by the geometric model. Moreover, the fact that the data lie *above* the point at which the line of the same colour cuts the ordinate indicates that the extent of perceived shearing was greater than the 2-D relative dot motion of the display. In contrast with the results of Experiment 11, there was no depth over-estimation in the 0.05 cpd surface.

Figs. 5.33 and 5.34 show the relative motion values which would theoretically be required to produce the perceived amounts of shearing and depth in each surface. The discrepancy between the theoretical and physical amplitudes is approximately the same for both corrugation frequencies. The absolute sizes of the discrepancies are generally smaller than those observed in the two previous experiments. As the physical amplitude increases, the discrepancy decreases somewhat, such that at the highest amplitude the theoretical magnitude is less than twice the physical magnitude.

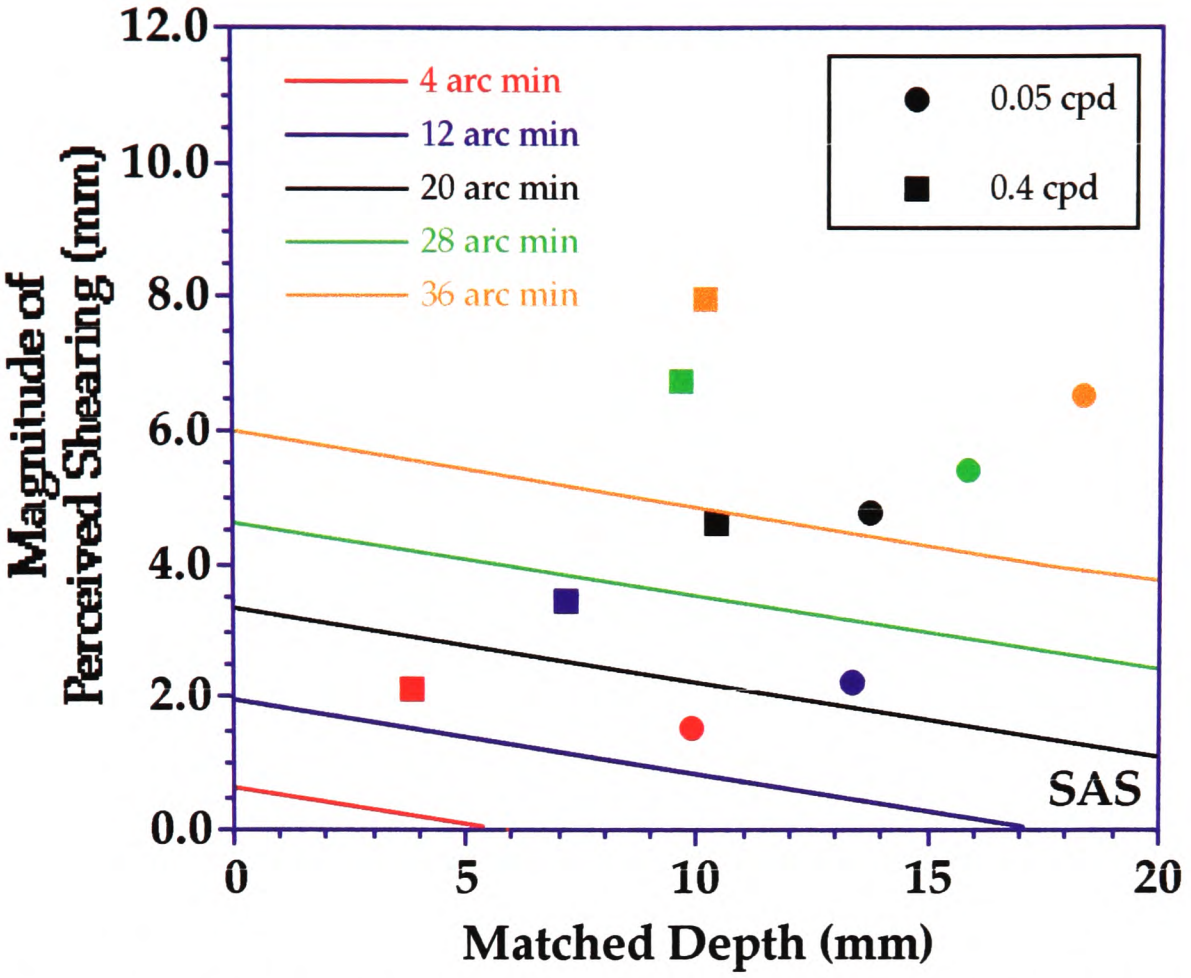


Fig. 5.31: Matched Shearing versus Matched Depth for two corrugation frequencies, for Subject SAS. The lines show the shearing/depth combinations predicted by the geometric model for the five different relative motion amplitudes. Data should theoretically fall on the line of the same colour.

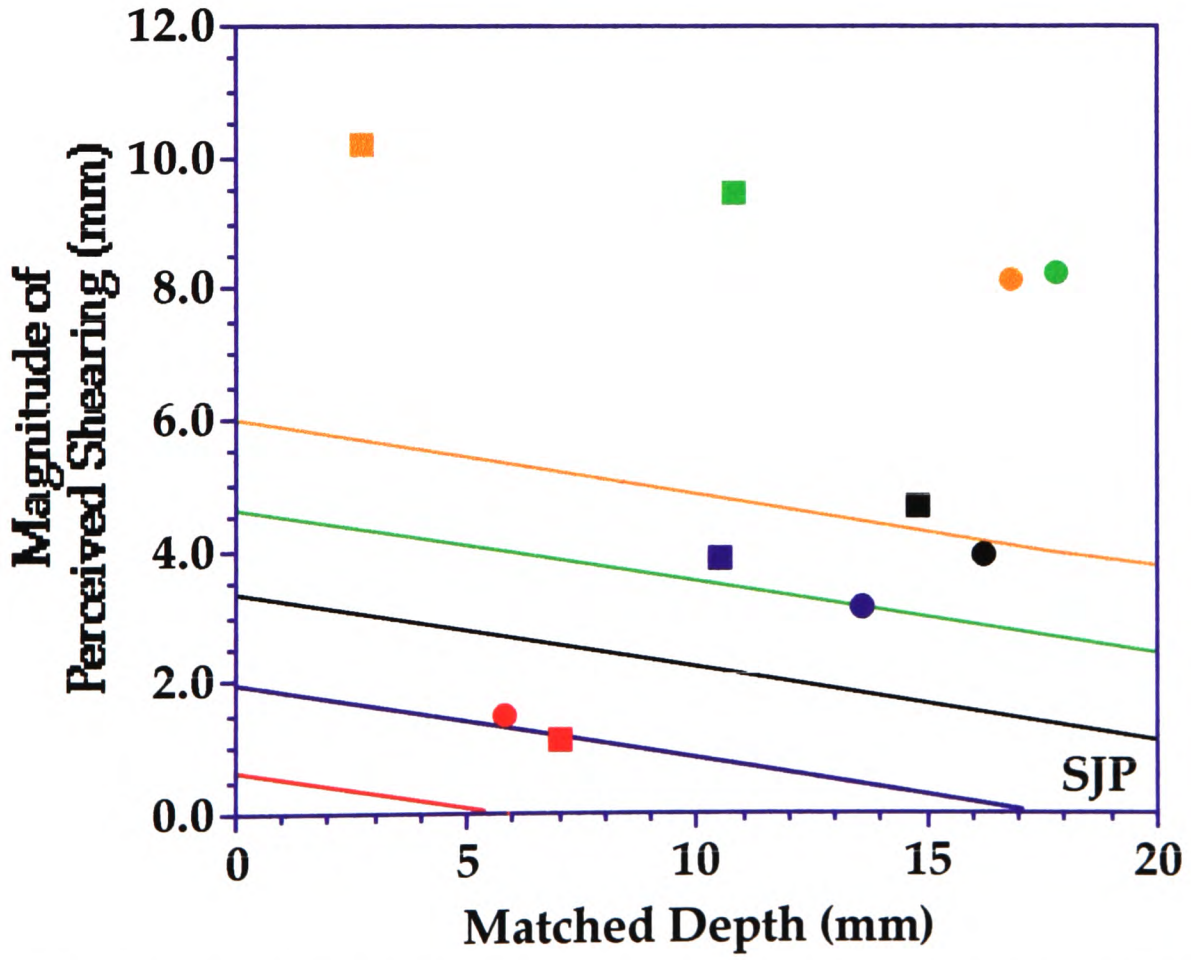


Fig. 5.32: Matched Shearing versus Matched Depth for two corrugation frequencies, for Subject SJP. The lines show the shearing/depth combinations predicted by the geometric model for the five different relative motion amplitudes. Data should theoretically fall on the line of the same colour.

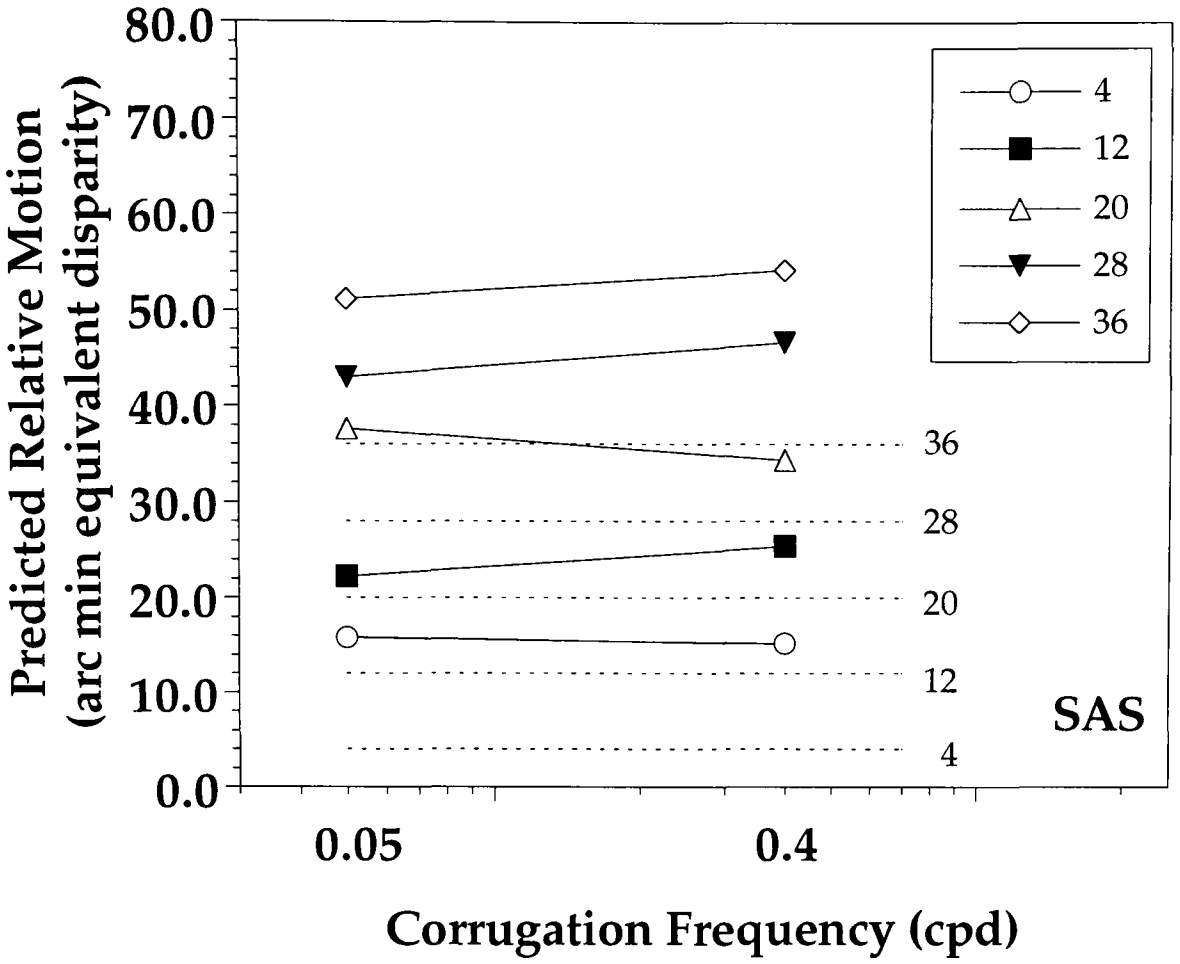


Fig 5.33: The relative motion values which would theoretically be required to produce the perceived amounts of shearing and depth in each surface, for Subject SAS. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

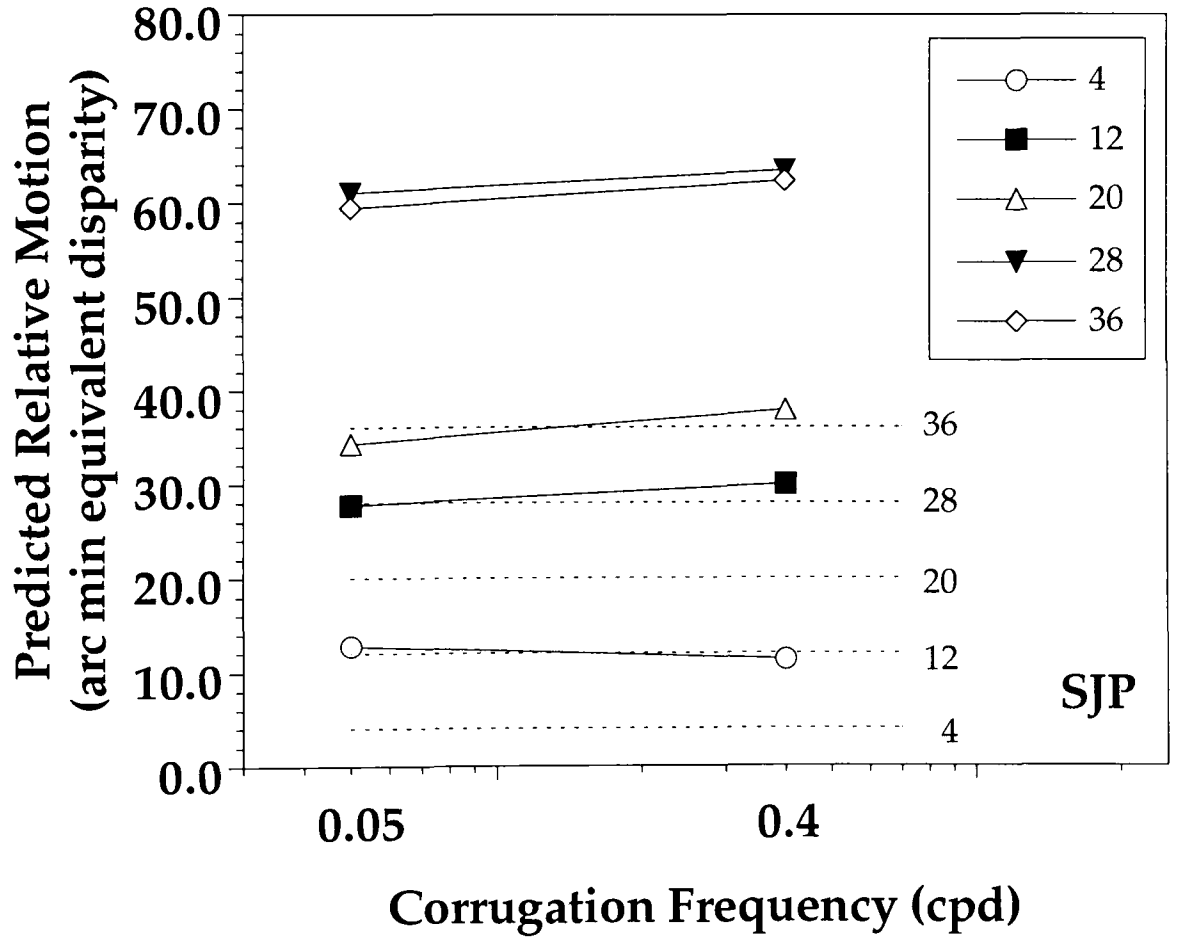


Fig 5.34: The relative motion values which would theoretically be required to produce the perceived amounts of shearing and depth in each surface, for Subject SJP. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

5.5.5 Discussion

These results are qualitatively consistent with a trade-off relationship between depth and shearing. As in the case of sinusoidal surfaces, the relative motion amplitudes theoretically required to produce the observed magnitudes of depth and shearing are greater than the physical amplitudes of the stimuli. The mean amounts of depth and shearing perceived in the 12 arc min equivalent disparity surface of Ono and Steinbach's (1990) study were approximately 9 mm and 1 mm respectively. These figures accord quite closely with those obtained for the 0.05 cpd surface of the current experiment, especially since it would be predicted that slightly less depth should be perceived in the 0.1 cpd surface of Ono and Steinbach.

5.6 Experiment 14: Matched Shearing and Matched Depth in Square Wave surfaces as a function of Stimulus Size

5.6.1 Stimuli

Pilot experiments established that the effect of size on the perceived amounts of depth and shearing was small, therefore stimulus size was either 2.5° or 20°. Corrugation frequency was 0.4 cpd.

5.6.2 Procedure

This was identical to the previous experiment.

5.6.3 Results

Figs. 5.35 and 5.37 show matched depth as a function of relative motion for the two stimulus sizes, for Subjects SAS and SJP respectively. A greater amount of depth was perceived in the 20° surface. The difference between the perceived depth magnitudes was about 1 - 5 mm. Again, Subject SJP perceived a decreasing amount of depth as the relative motion amplitude increased beyond 20 arc min equivalent disparity. This subject perceived no depth in the smaller stimulus at the highest amplitude. Figs. 5.36 and

5.38 depict the perceived magnitude of shearing as a function of stimulus size, for Subjects SAS and SJP respectively. A greater magnitude of shearing was perceived in the 20° surface than in the 2.5° surface.

Figs. 5.39 and 5.40 show the data plotted with respect to the depth/shearing combinations predicted by the model. Again, the coloured legend represents the relative motion amplitudes (arc min equivalent disparity) which produce each theoretical depth/shearing curve. The right-hand legend shows stimulus size. Symbols of identical colour arose from surfaces containing identical relative motion amplitudes. The data should theoretically lie on the curve of the same colour. For Subject SAS, the combined amounts of depth and shearing perceived in the 2.5° stimulus lay close to the predicted amounts. For the 20° stimulus, and for both stimuli for Subject SJP, there was an over-estimation of the combined amounts of depth and shearing relative to those predicted by the geometric model. Once more for both subjects, the extent of perceived shearing was frequently greater than the extent of 2-D relative dot motion in the display.

Figs. 5.41 and 5.42 show the relative motion values which would theoretically be required to produce the perceived amounts of shearing and depth in each surface. As in Experiment 13, increases in the physical amplitude lead to decreases in the discrepancy between the physical and theoretical amplitude. Once again, the theoretical amplitudes are significantly lower than the corresponding amplitudes of Experiment 12. They are slightly lower for the 2.5° stimulus.

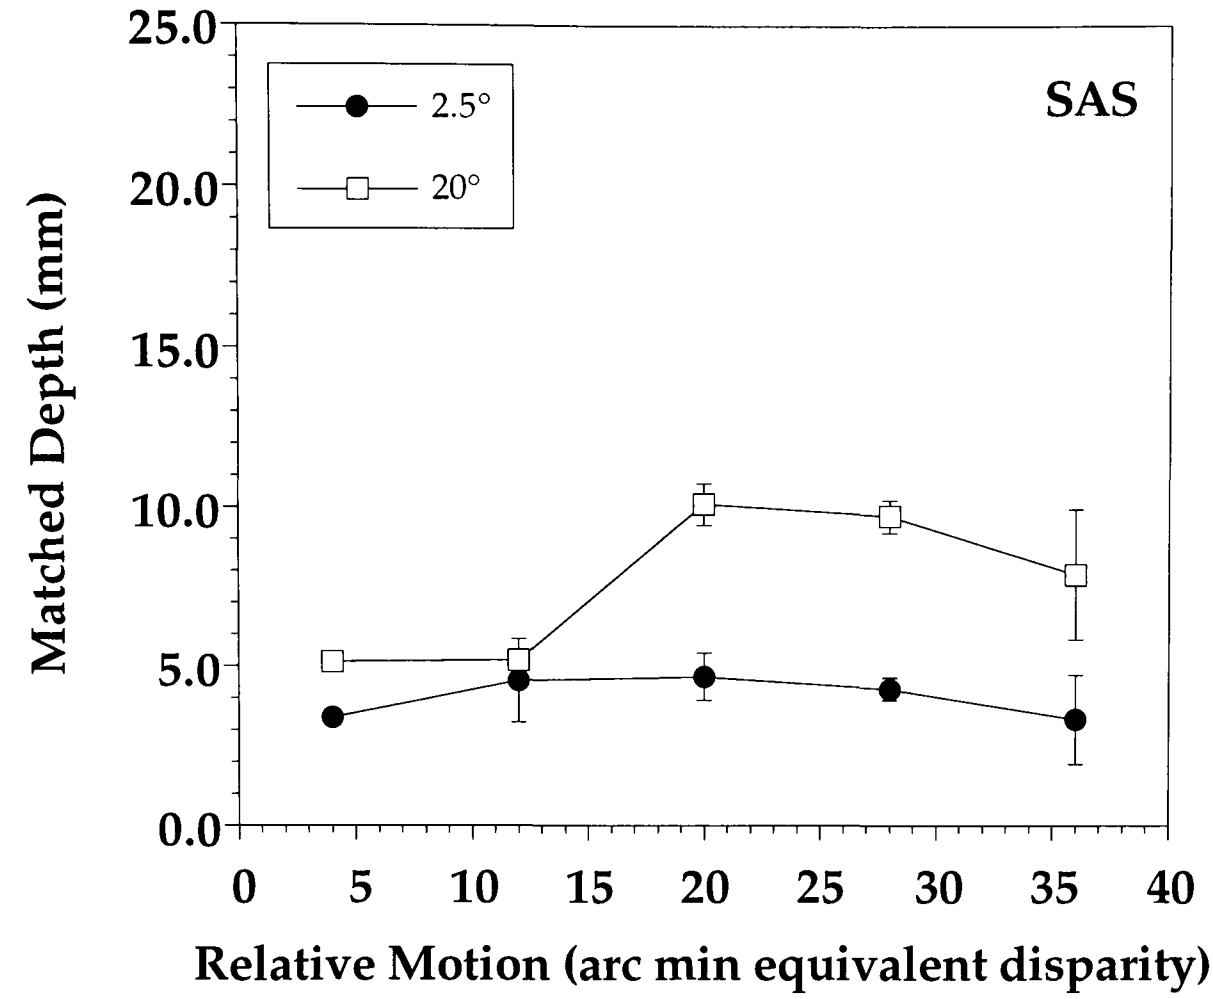


Fig. 5.35: Matched Depth as a function of relative motion amplitude, for Subject SAS.

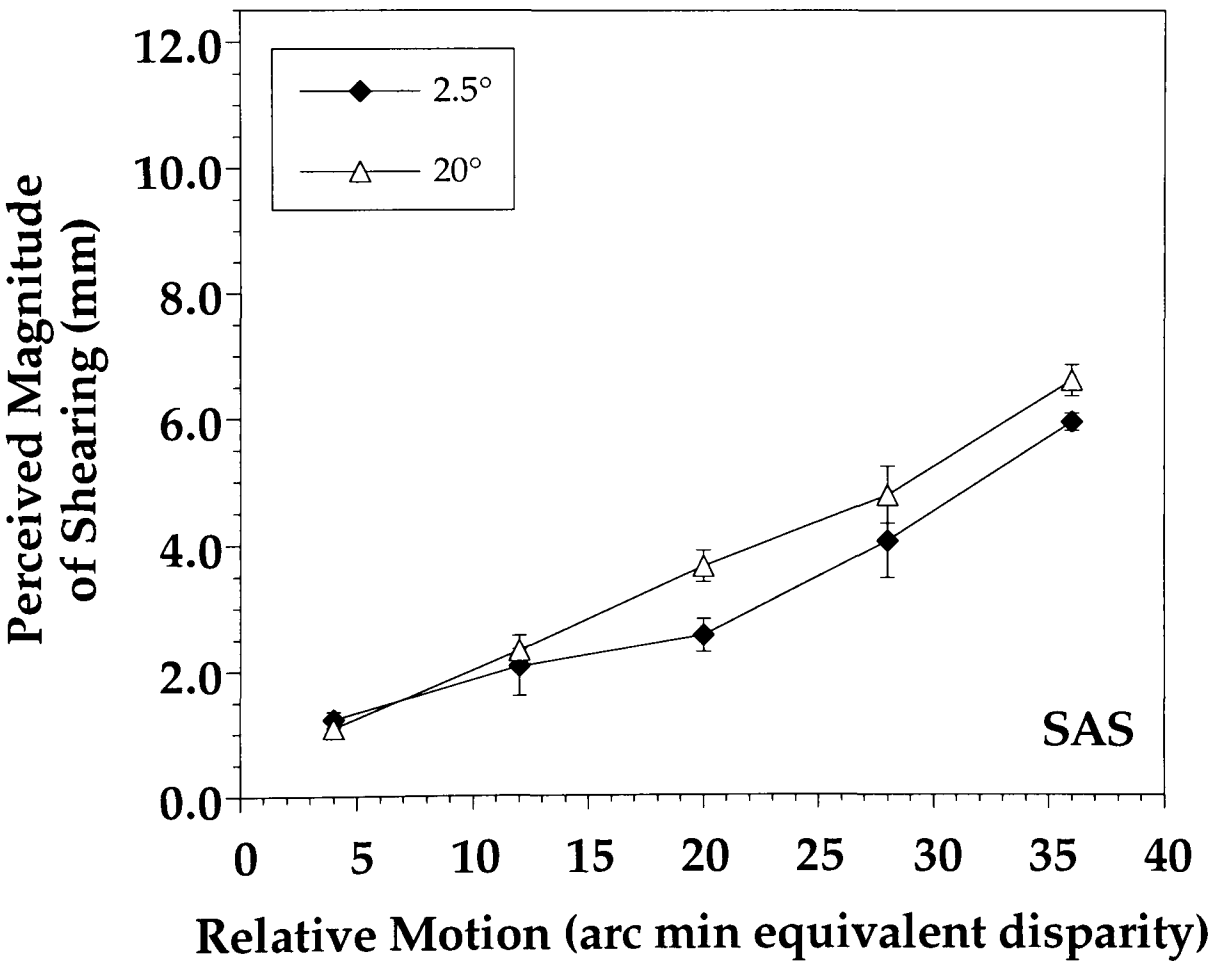


Fig. 5.36: Matched Shearing as a function of relative motion amplitude, for Subject SAS.

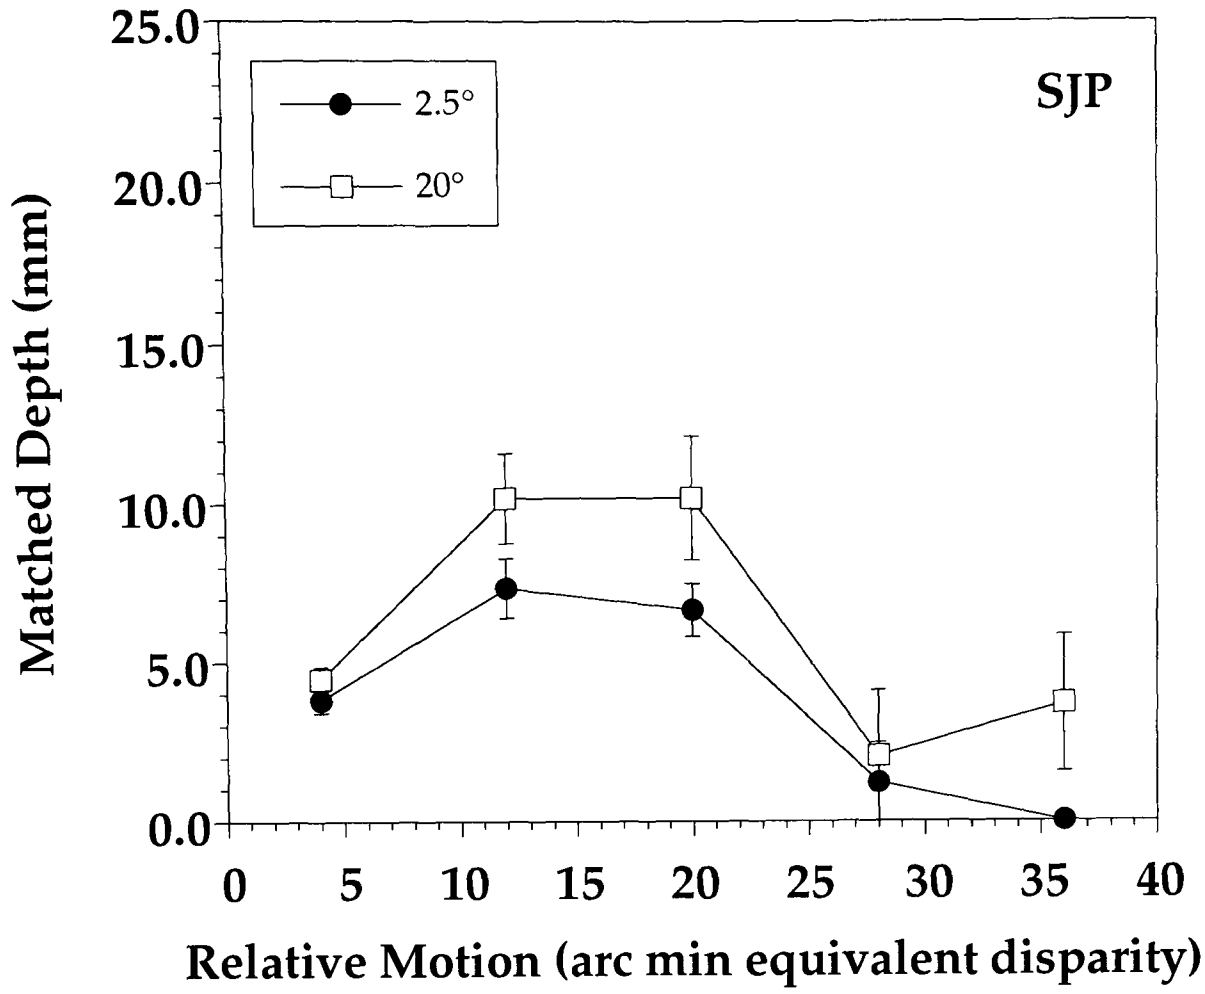


Fig. 5.37: Matched Depth as a function of relative motion amplitude, for Subject SJP.

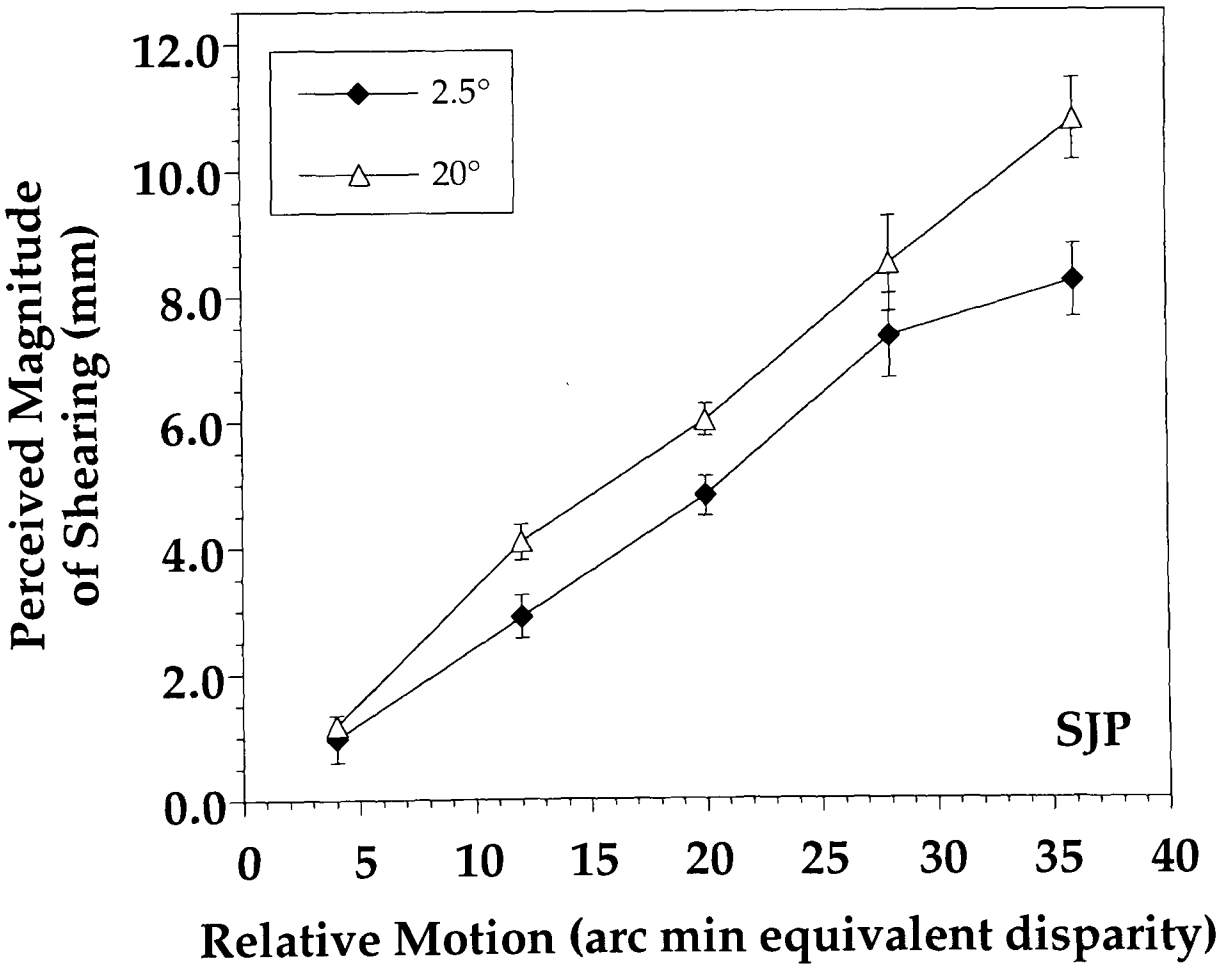


Fig. 5.38: Matched Shearing as a function of relative motion amplitude, for Subject SJP.

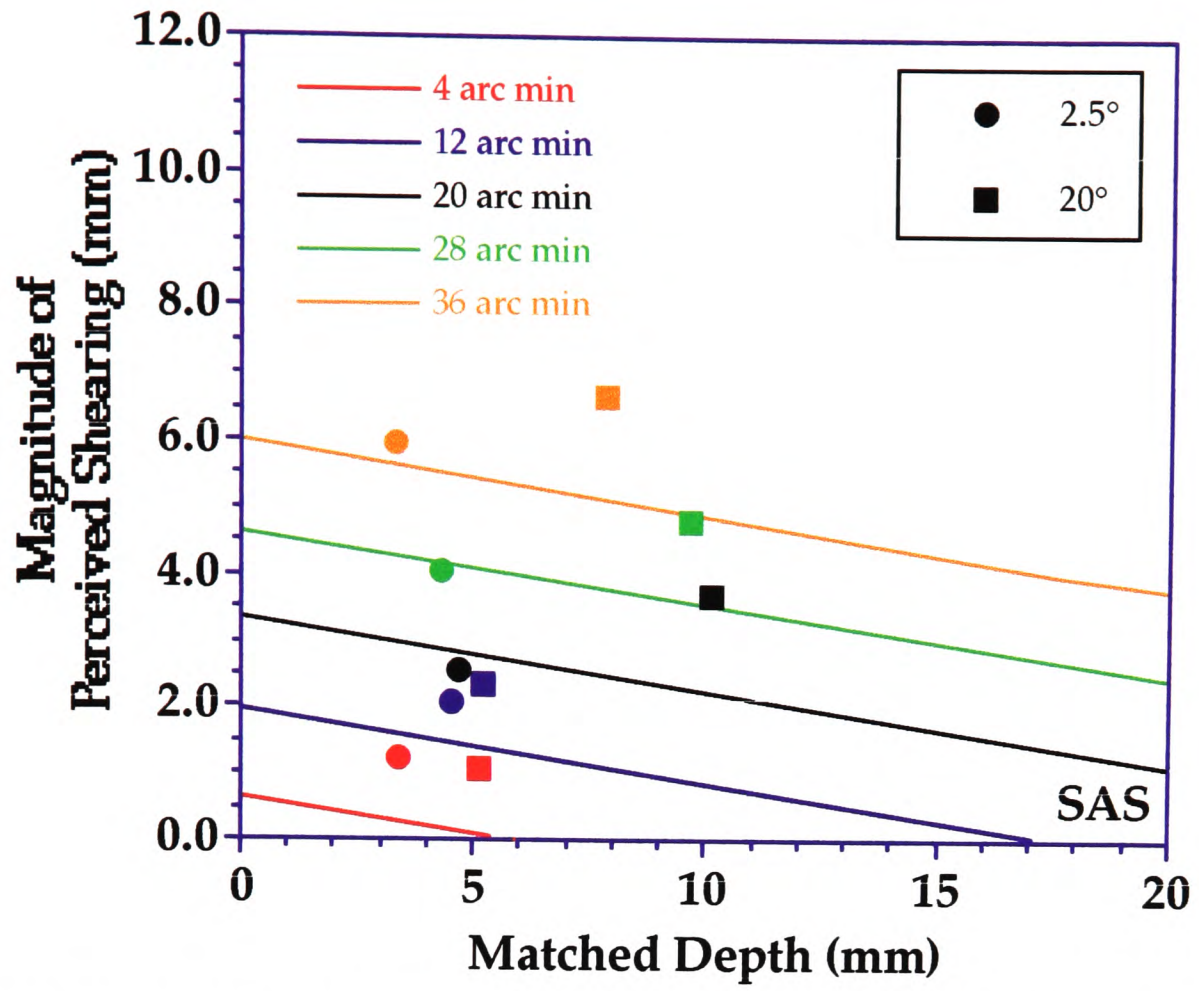


Fig. 5.39: Matched Shearing versus Matched Depth for two stimulus sizes, for Subject SAS. The lines show the shearing/depth combinations predicted by the geometric model for the five different relative motion amplitudes. Data should theoretically fall on the line of the same colour.

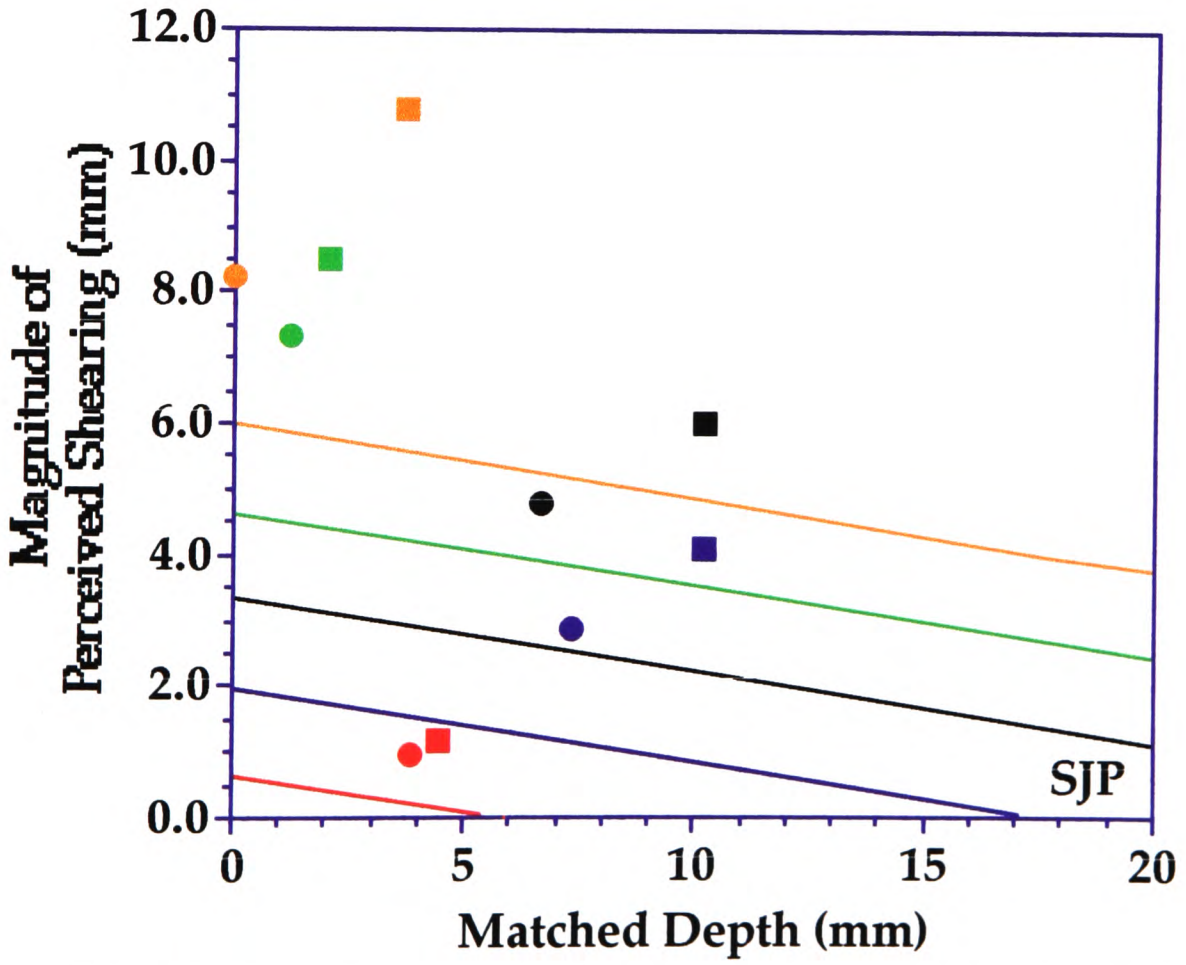


Fig. 5.40: Matched Shearing versus Matched Depth for two stimulus sizes, for Subject SJP. The lines show the shearing/depth combinations predicted by the geometric model for the five different relative motion amplitudes. Data should theoretically fall on the line of the same colour.

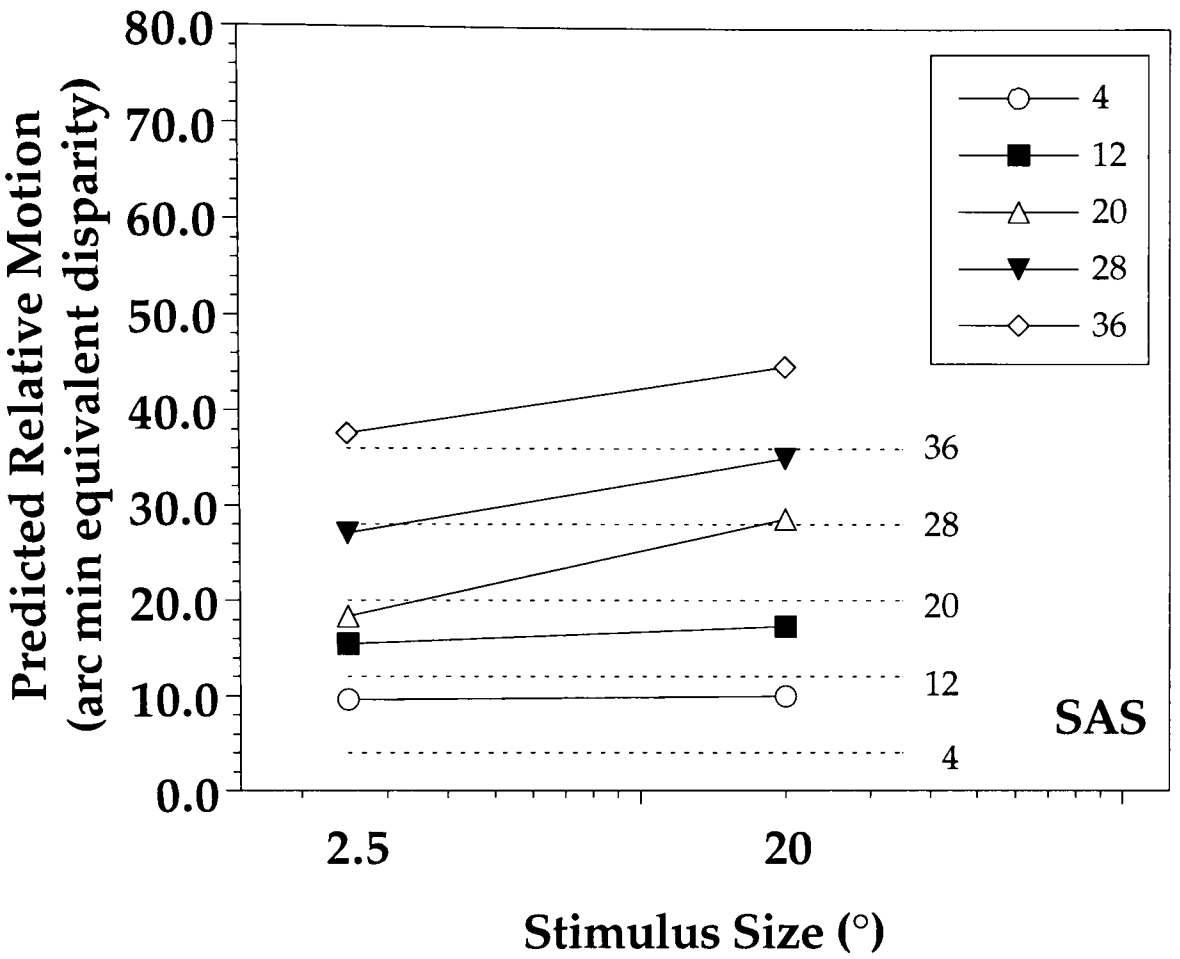


Fig 5.41: The relative motion values which would theoretically be required to produce the perceived amounts of shearing and depth in each surface, for Subject SAS. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

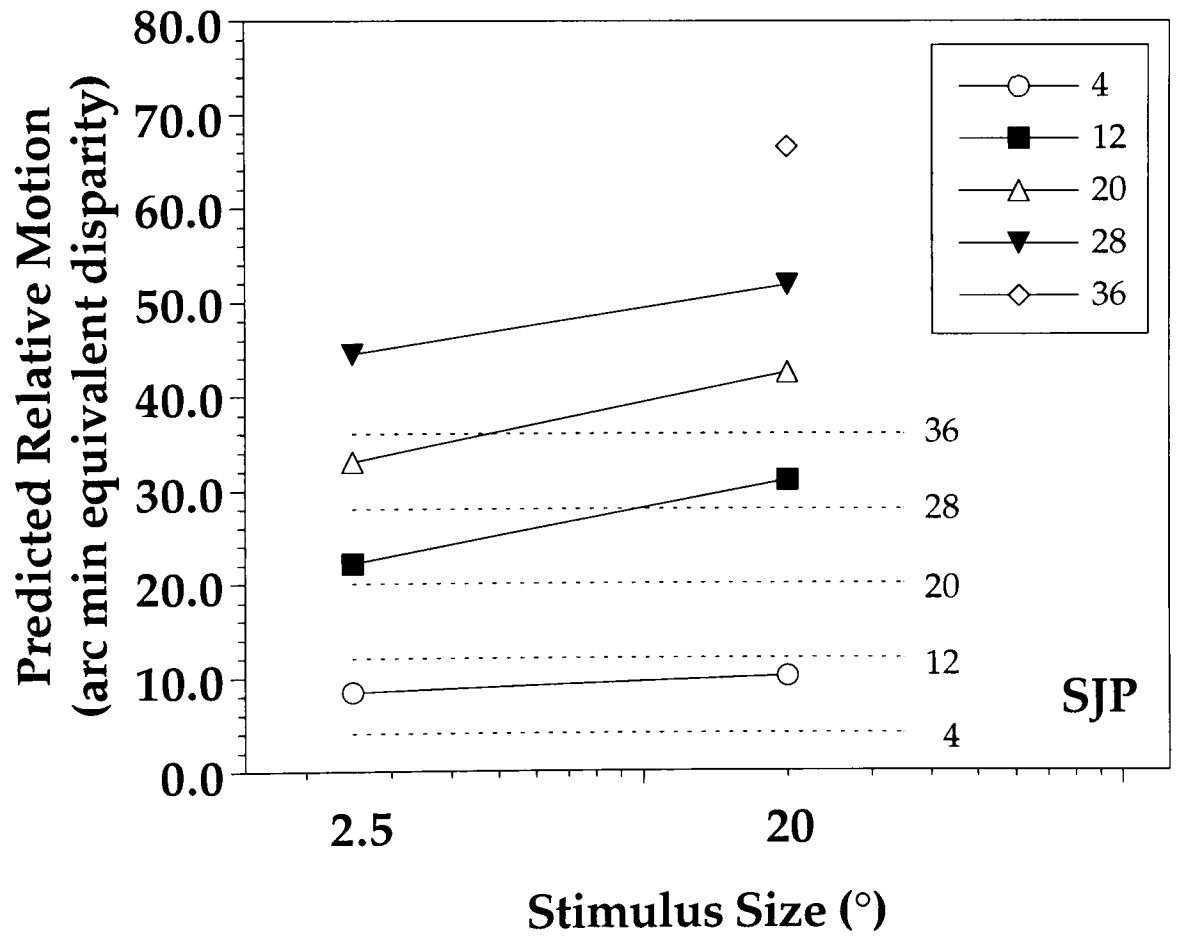


Fig 5.42: The relative motion values which would theoretically be required to produce the perceived amounts of shearing and depth in each surface, for Subject SJP. The legend shows the physical relative motion amplitudes (arc min equivalent disparity) of the surfaces. These amplitudes are represented by the dashed lines.

5.6.4 Discussion

The perceived magnitudes of depth and shearing were not related by the predicted trade-off relationship. Instead, the larger stimulus produced greater perceived extents of both attributes. Once again, however, it was found that the combined extents of depth and shearing were greater than those predicted by the geometric model. Interestingly, increases in relative motion amplitude beyond about 20 arc min equivalent disparity did not result in further increases in perceived depth for one subject, and resulted in an increasing loss of depth for subject SJP.

5.7 Experiment 15: Perceived Rotation and Perceived Depth in Proximally Identical Stimuli as a function of Viewing Distance

5.7.1 Introduction

Ono, Rivest and Ono (1986) investigated depth perception as a function of motion parallax and absolute distance information. Observers were asked to indicate the amount of perceived depth in sinusoidally corrugated parallax surfaces at viewing distances of either 40, 80, 160 or 320 cm, and to indicate whether or not the surface appeared to rotate. Three levels of distal parallax were used. It was found that as viewing distance increased, the amount of depth was increasingly under-estimated. Moreover, the number of observers perceiving the parallax surface to rotate increased also. The observed trade-off between depth and rotation led to the claim that the 'limit of effective motion parallax cannot be specified simply in terms of proximal parallax when absolute-distance information is available to the visual system' (p. 335). The term 'effective motion parallax' was used by these authors to refer to the situation in which relative motion was interpreted solely as depth, with no rotation. Using the terminology of this thesis, the 'limit of effective parallax' refers to the Transition Point. Experiment 4 eliminated the confounding variables of Ono et al.'s (1986) study (see Sections 3.2.3 and 3.3), and it was concluded that the Transition Point falls with increasing viewing distance.

Although Ono et al.'s (1986) study showed that for a given amount of proximal parallax more observers perceived rotation as viewing distance increased, quantitative measurements of perceived rotation were not taken. The authors' claim that rotation was perceived at the expense of depth is investigated more rigorously in the current experiment.

5.7.2 Stimuli

There were two sets of stimuli. The stimuli within each set were proximally identical in terms of corrugation frequency, angle subtended by each display element, display size and amplitude of horizontal relative motion.

The first set consisted of stimuli which subtended 5° and which contained relative motion amplitudes of 2 - 10 arc min equivalent disparity. These stimuli were viewed from distances of 57, 114 and 228 cm. The second set consisted of stimuli which subtended 10° and which contained relative motion amplitudes of 4 - 20 arc min equivalent disparity. These stimuli were viewed from distances of 57 and 114 cm.

All stimuli had a corrugation frequency of 0.4 cpd.

5.7.3 Procedure

Trials were blocked in groups of ten, corresponding to two repetitions of each of the relative motion amplitudes. Subjects completed three experimental sessions, each consisting of three blocks (one at each viewing distance). The order of blocks within each session was randomised. For Subject SAS, each trial consisted of a Depth Match and a Rotation Match. In half the total number of trials the depth judgement was made first and the rotation judgement was made second. In the other half, this order was reversed. For the two other subjects, Depth Match data were collected on a later occasion than the Rotation Match data. Subject SAS repeated the Depth Matching at the time at which it was carried out by the remaining subjects; her two sets of Depth Match data were similar.

5.7.4 Results

Figs. 5.43 - 5.48 show Matched Rotation and Matched Depth for stimuli which subtended 5° and which contained relative motion amplitudes of 2 - 10 arc min equivalent disparity. For all subjects, there was an increase in the magnitudes of perceived rotation and depth as viewing distance increased. Subject SAS perceived lower extents of rotation than the other two subjects.

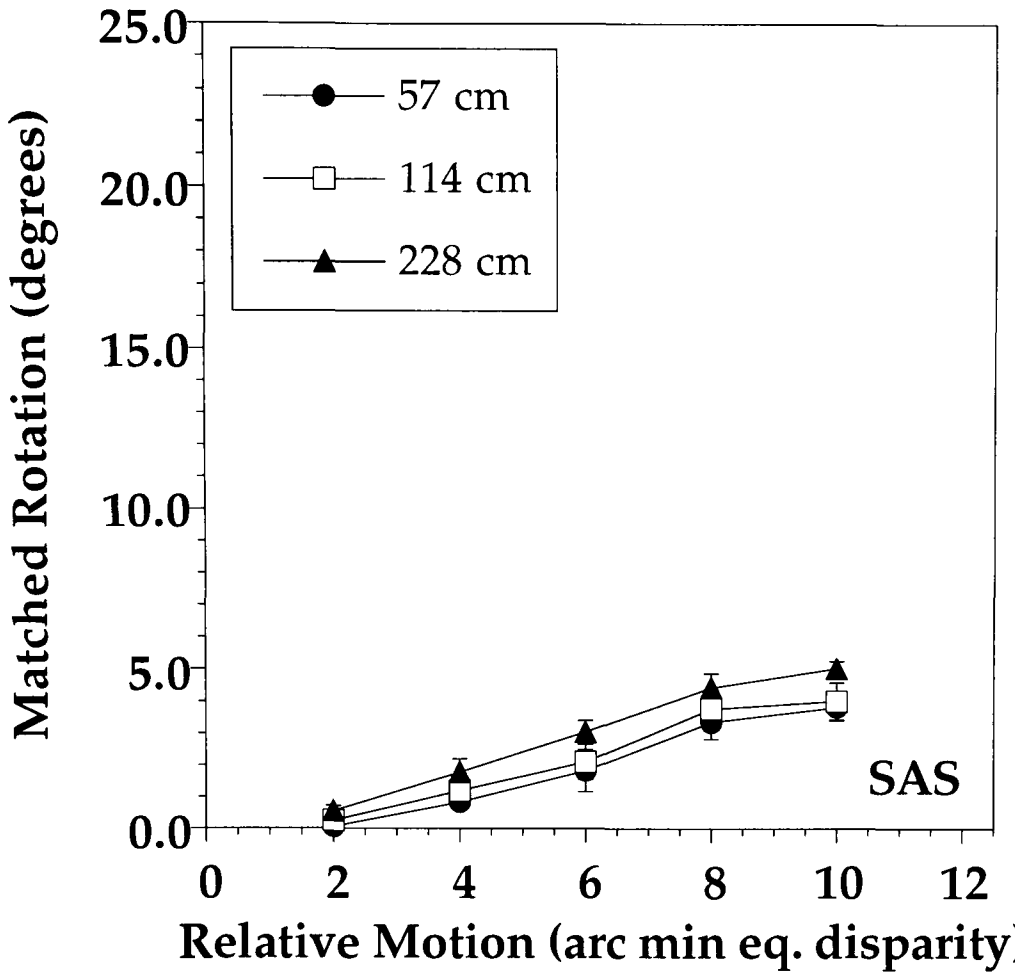


Fig 5.43: Matched Rotation as a function of relative motion amplitude for proximally identical stimuli of size 5° at three viewing distances, for Subject SAS.

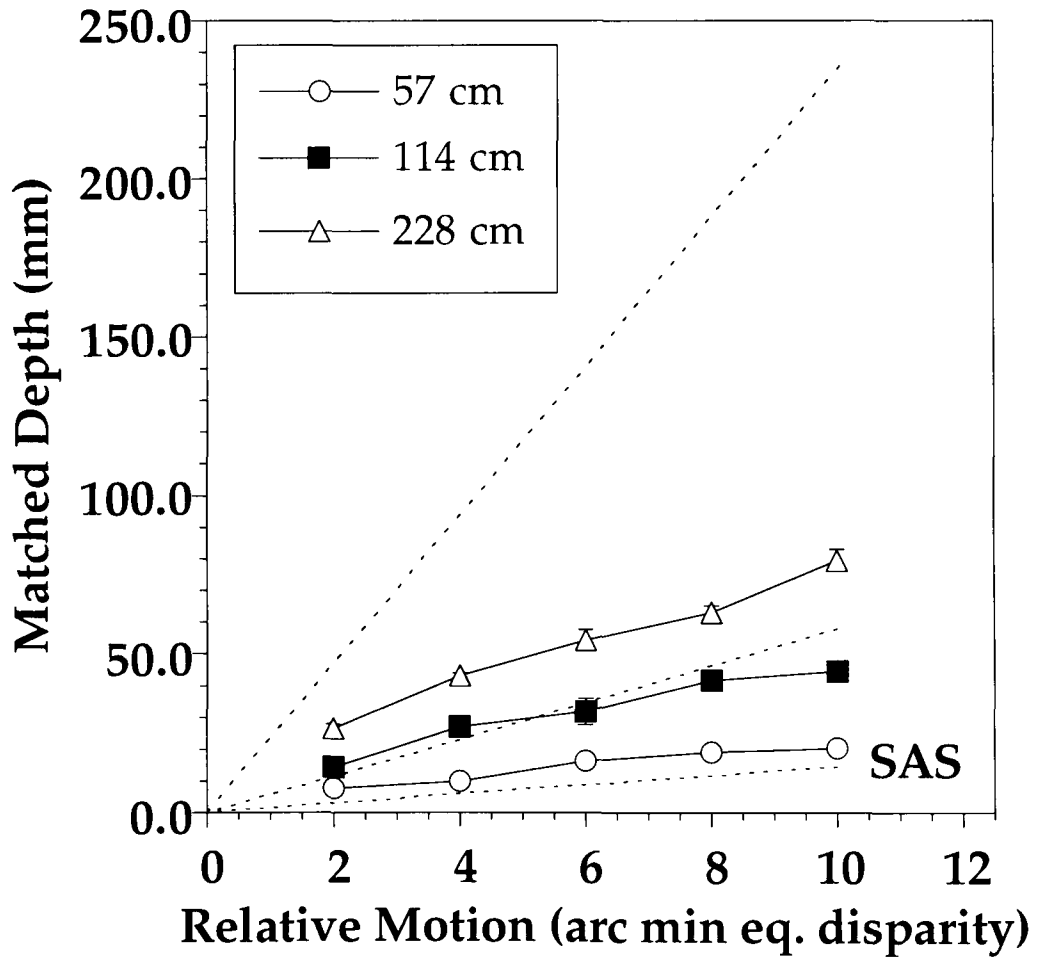


Fig 5.44: Matched Depth as a function of relative motion amplitude for proximally identical stimuli of size 5° at three viewing distances, for Subject SAS. The dashed lines represent the simulated depth for viewing distances of 228 cm (steepest slope), 114 cm and 57 cm.

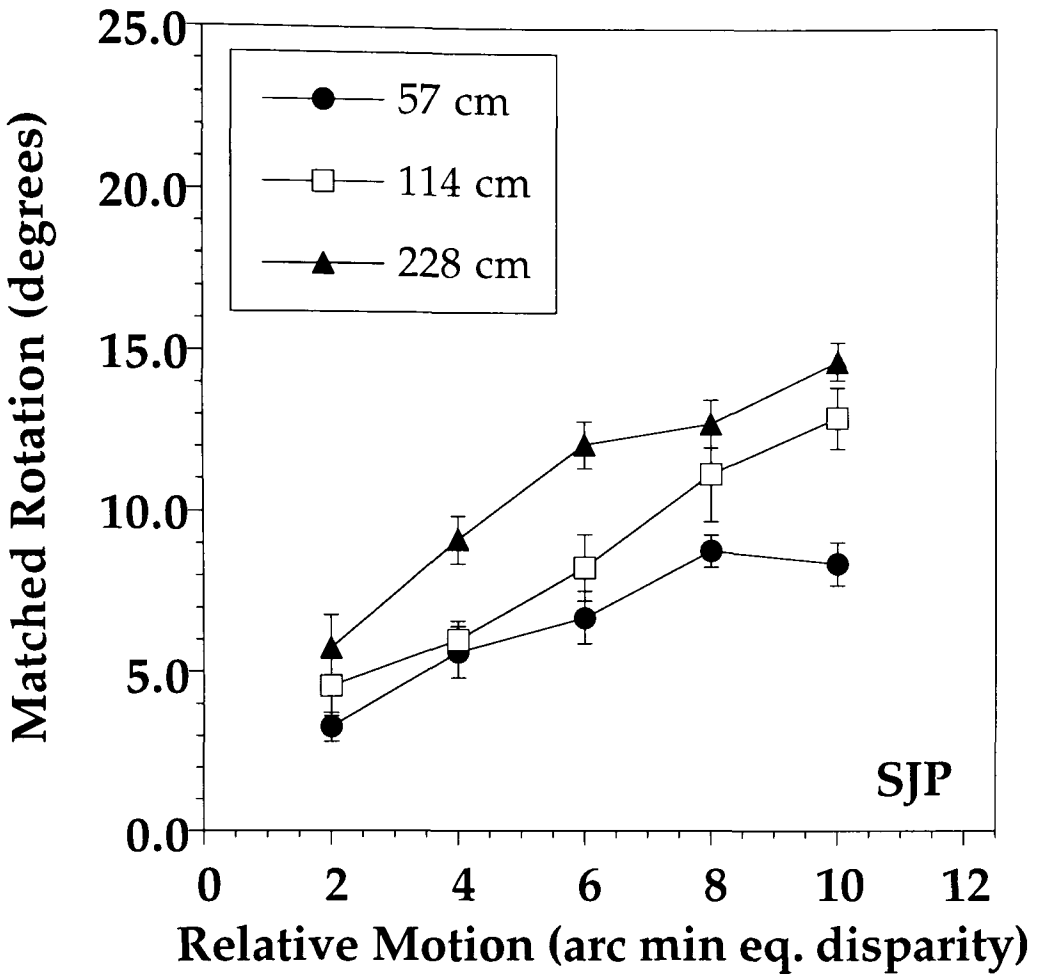


Fig 5.45: Matched Rotation as a function of relative motion amplitude for proximally identical stimuli of size 5° at three viewing distances, for Subject SJP.

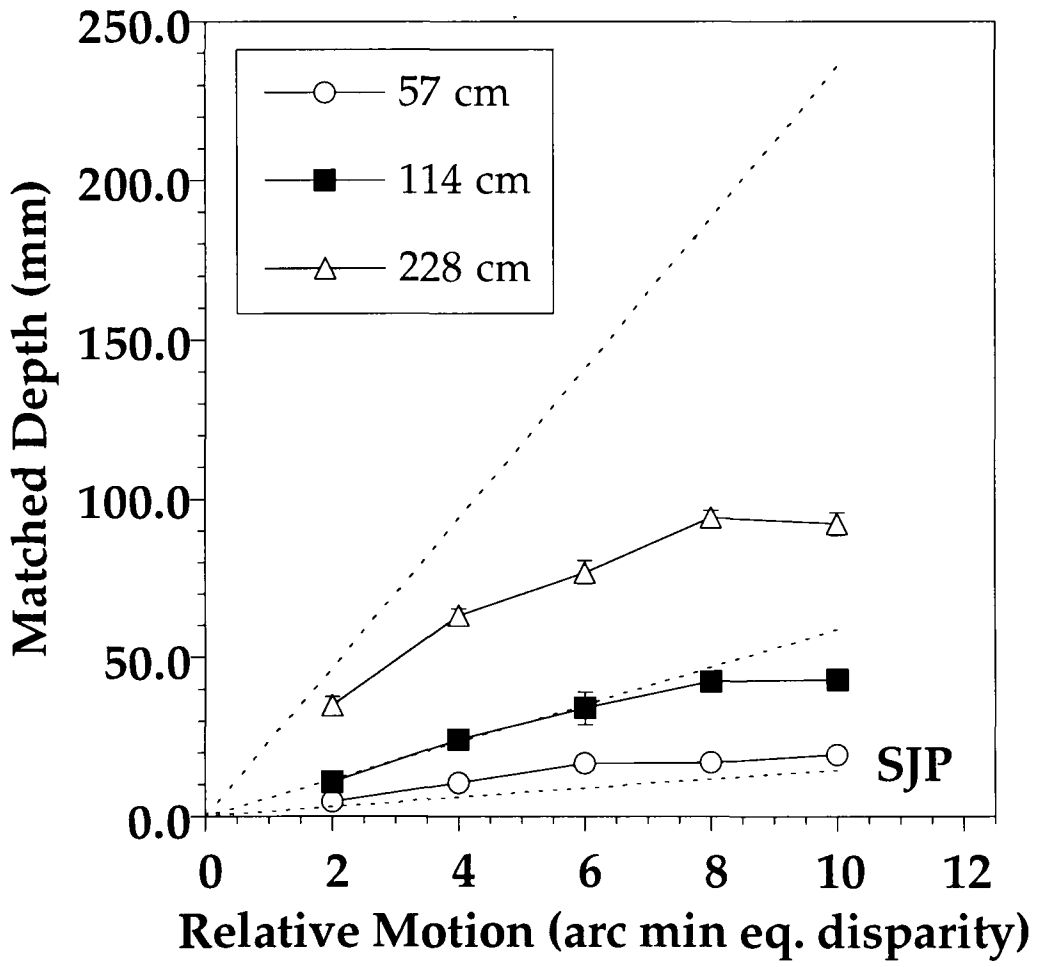


Fig 5.46: Matched Depth as a function of relative motion amplitude for proximally identical stimuli of size 5° at three viewing distances, for Subject SJP. The dashed lines represent the simulated depth for viewing distances of 228 cm (steepest slope), 114 cm and 57 cm.

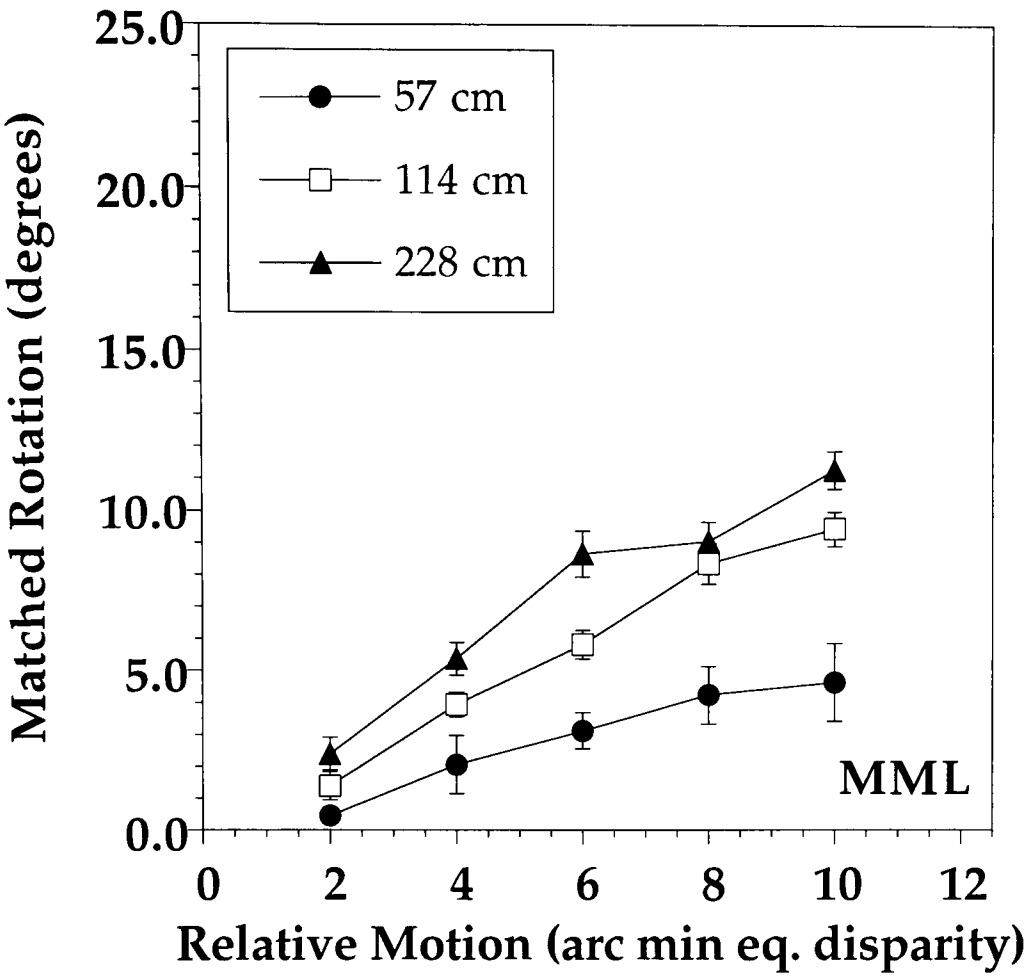


Fig 5.47: Matched Rotation as a function of relative motion amplitude for proximally identical stimuli of size 5° at three viewing distances, for Subject MML.

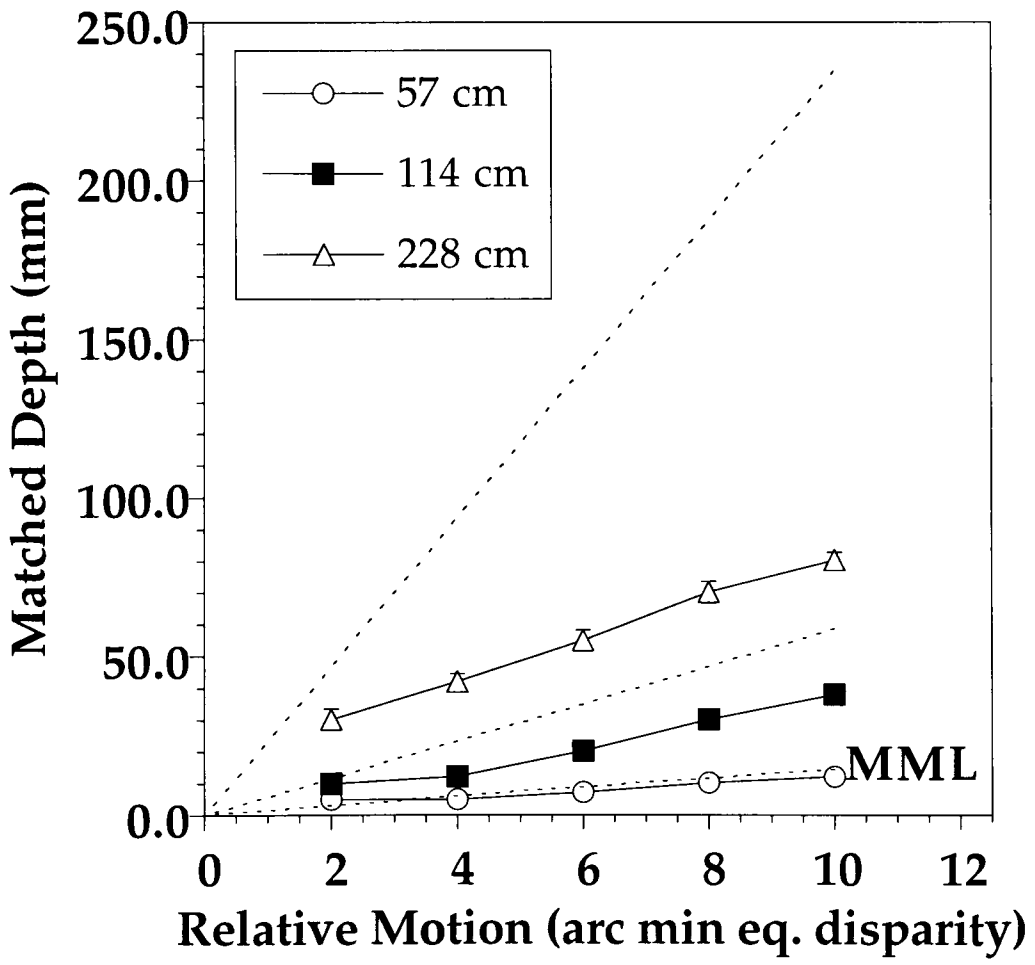


Fig 5.48: Matched Depth as a function of relative motion amplitude for proximally identical stimuli of size 5° at three viewing distances, for Subject MML. The dashed lines represent the simulated depth for viewing distances of 228 cm (steepest slope), 114 cm and 57 cm.

Figs. 5.49 - 5.54 show Matched Rotation and Matched Depth for stimuli which subtended 10° and which contained relative motion amplitudes of 4 - 20 arc min equivalent disparity. Again, there was an increase in the magnitudes of perceived depth as viewing distance increased. The rotation functions at the two distances were not well separated for Subject SAS. For the other two subjects, a greater extent of rotation was perceived at the further viewing distance.

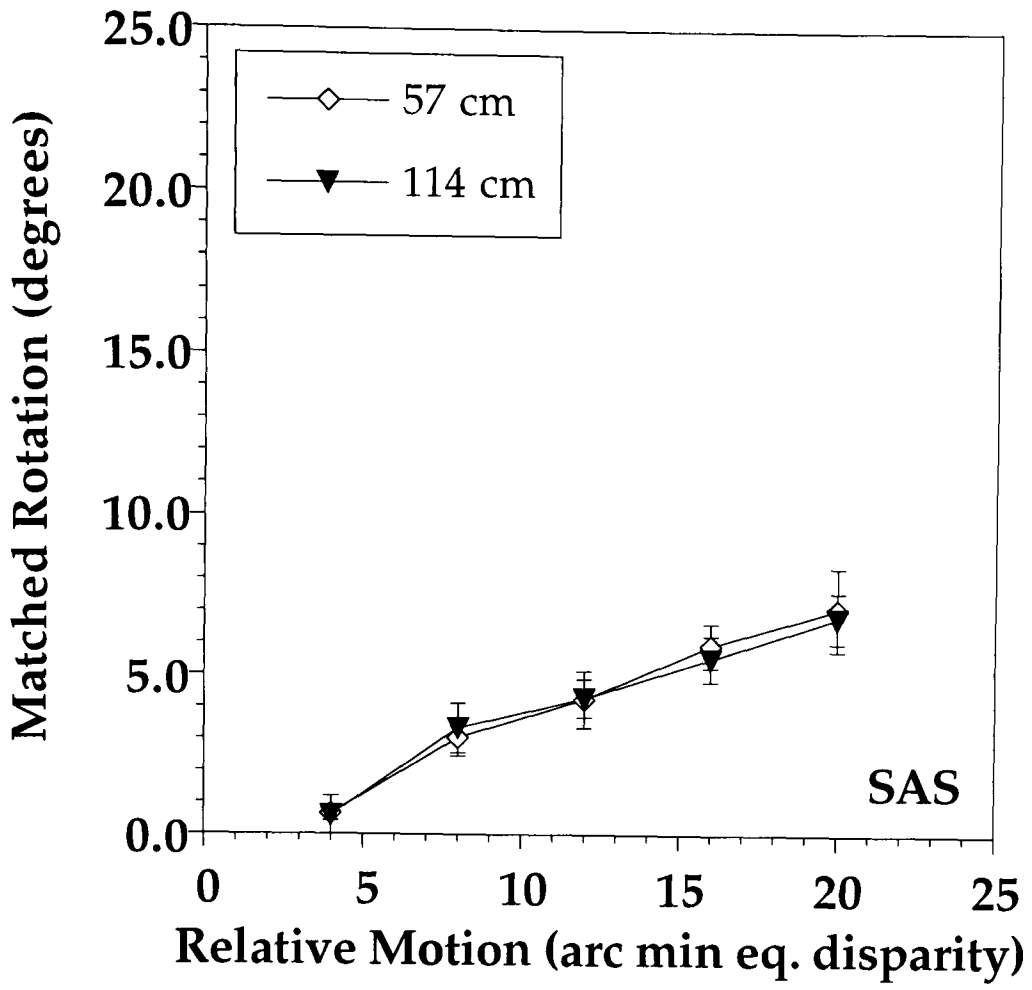


Fig. 5.49: Matched Rotation as a function of relative motion amplitude for proximally identical stimuli of size 10° at two viewing distances, for Subject SAS.

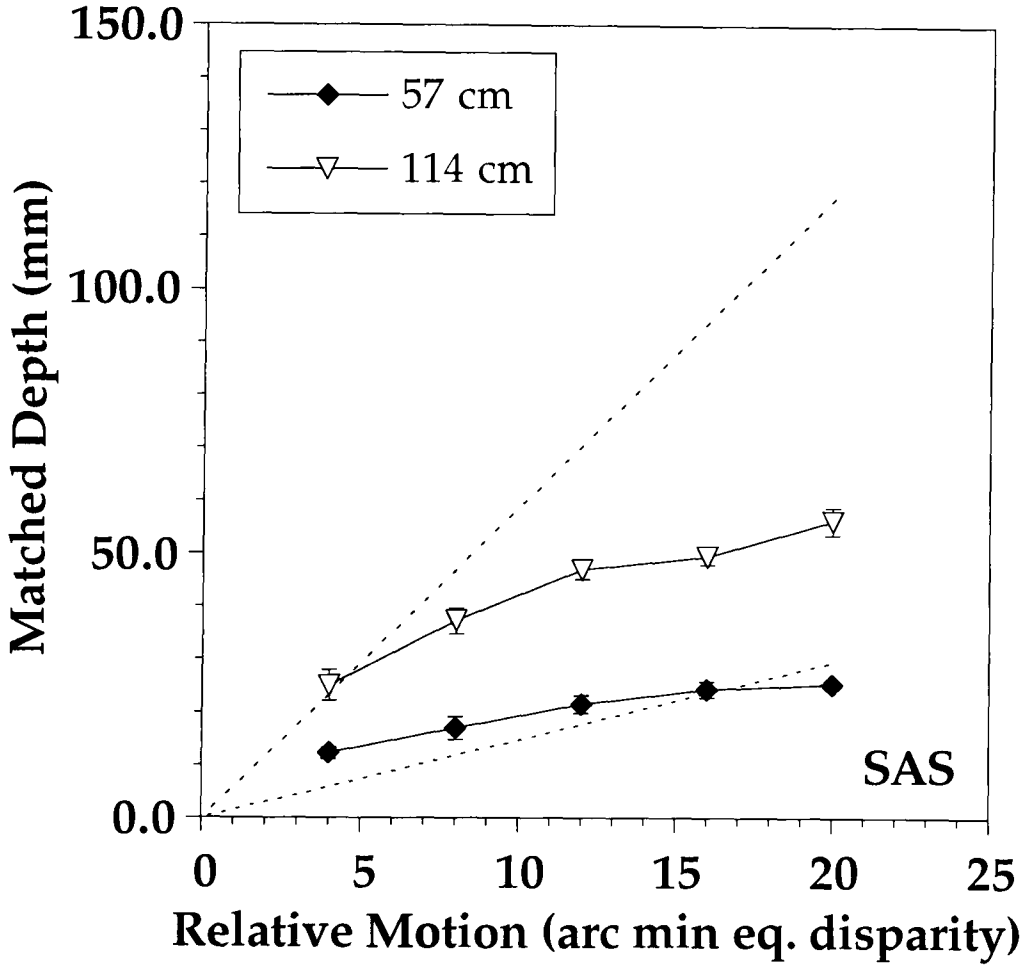


Fig. 5.50: Matched Depth as a function of relative motion amplitude for proximally identical stimuli of size 10° at two viewing distances, for Subject SAS. The dashed lines represent the simulated depth for viewing distances of 114 cm (steeper slope) and 57 cm.

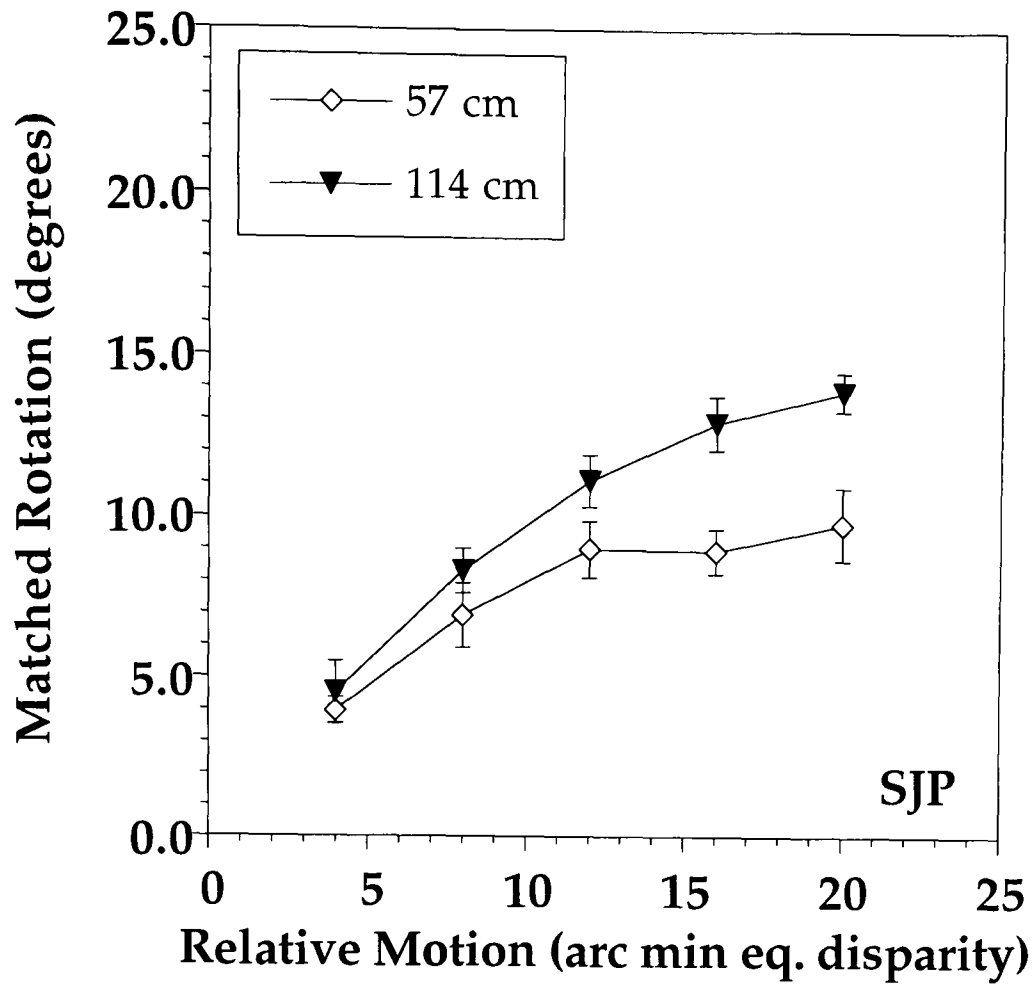


Fig. 5.51: Matched Rotation as a function of relative motion amplitude for proximally identical stimuli of size 10° at two viewing distances, for Subject SJP.

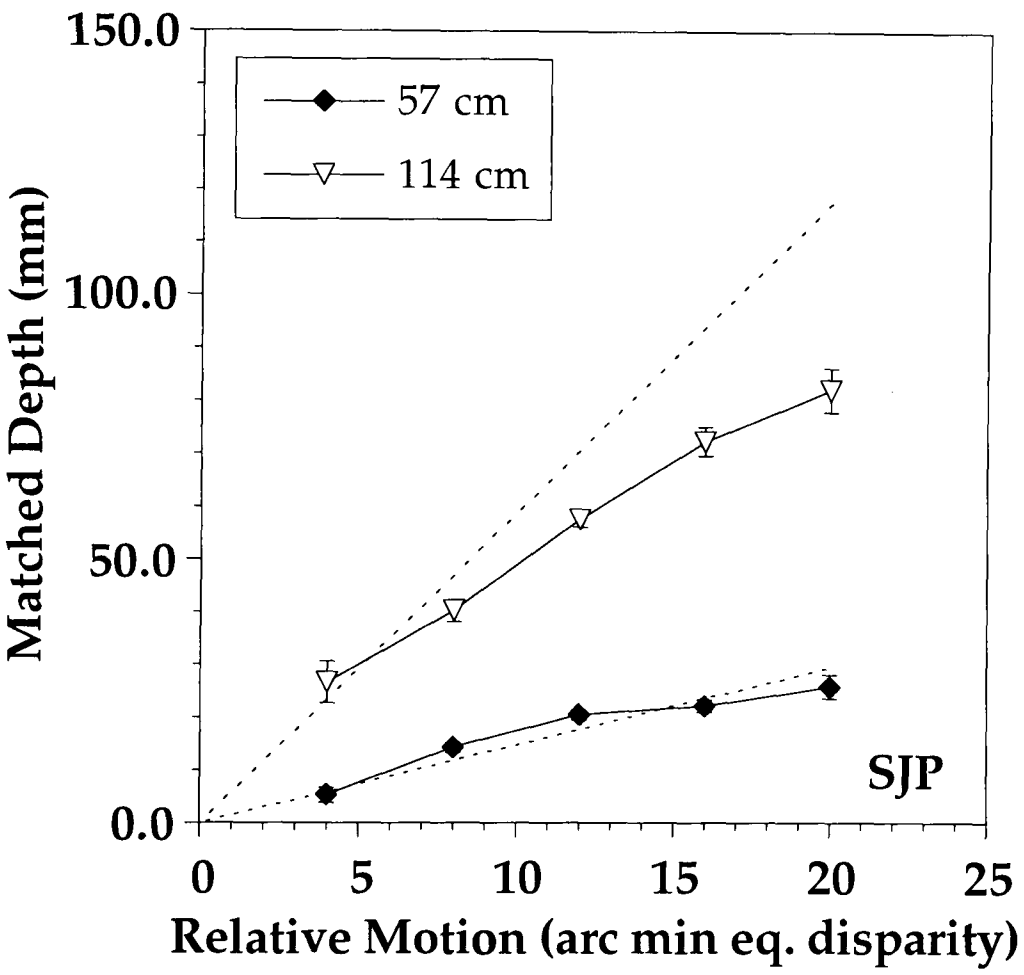


Fig. 5.52: Matched Depth as a function of relative motion amplitude for proximally identical stimuli of size 10° at two viewing distances, for Subject SJP. The dashed lines represent the simulated depth for viewing distances of 114 cm (steeper slope) and 57 cm.

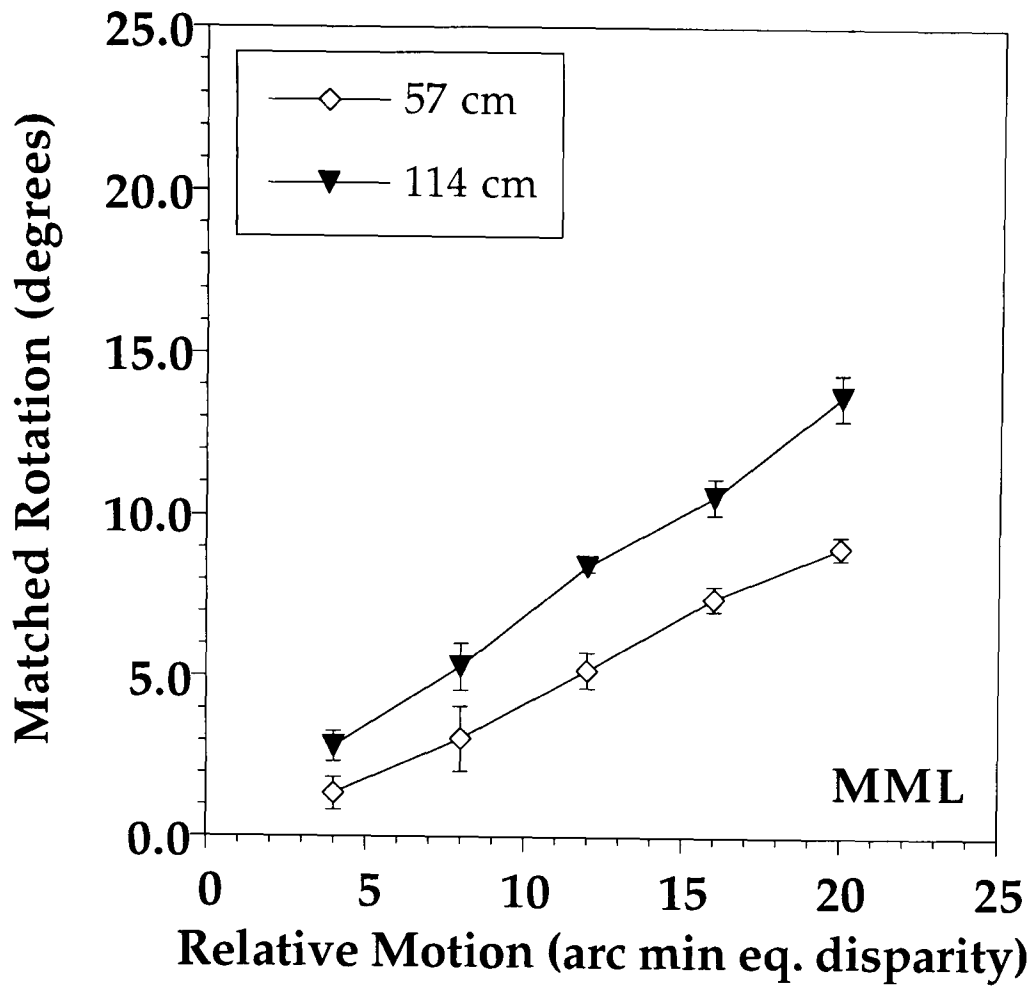


Fig. 5.53: Matched Rotation as a function of relative motion amplitude for proximally identical stimuli of size 10° at two viewing distances, for Subject MML.

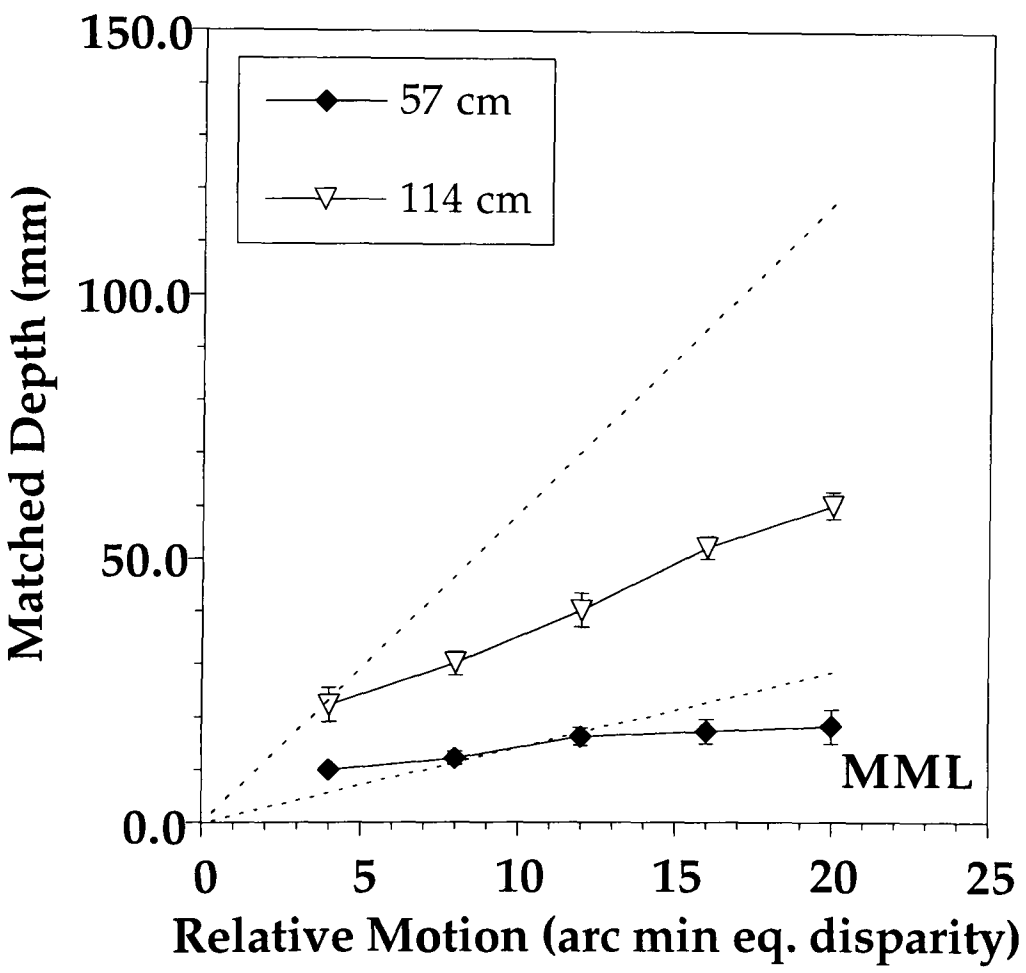
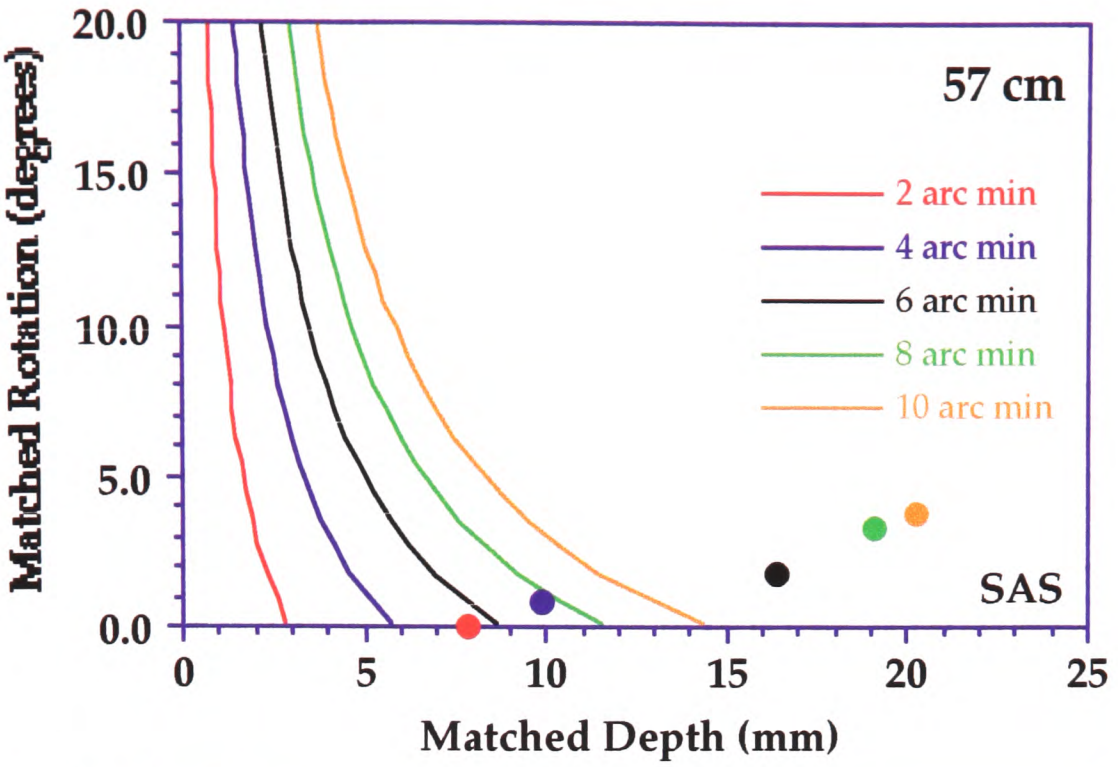
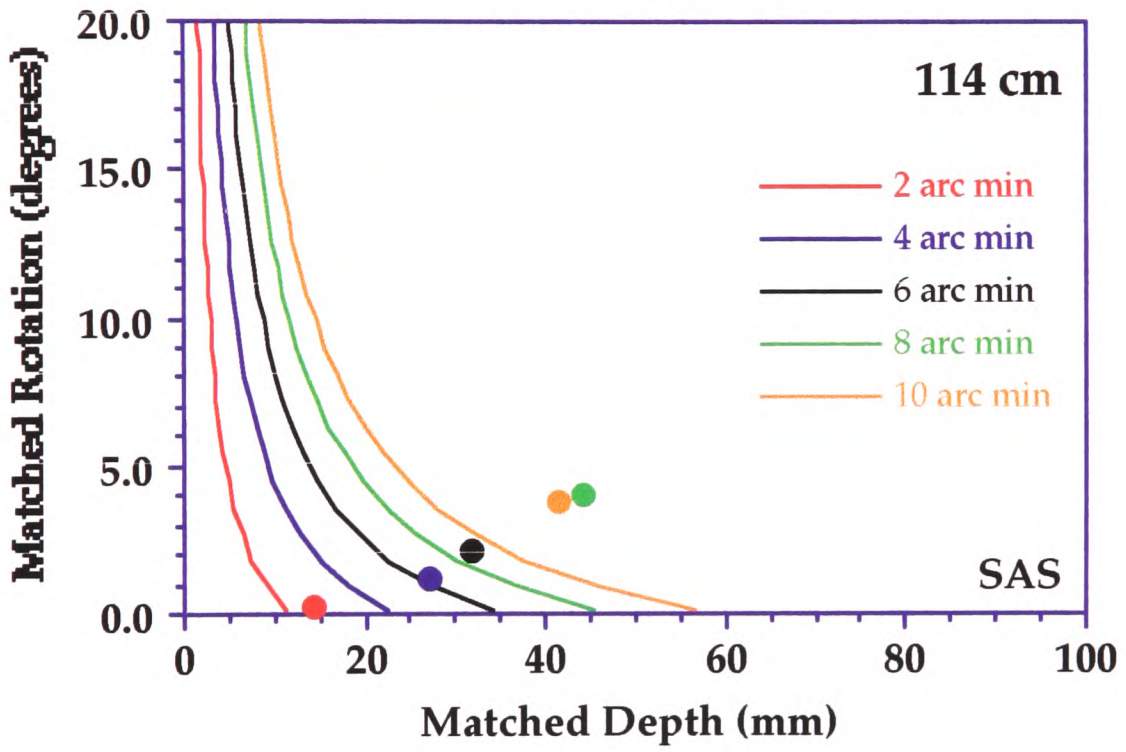
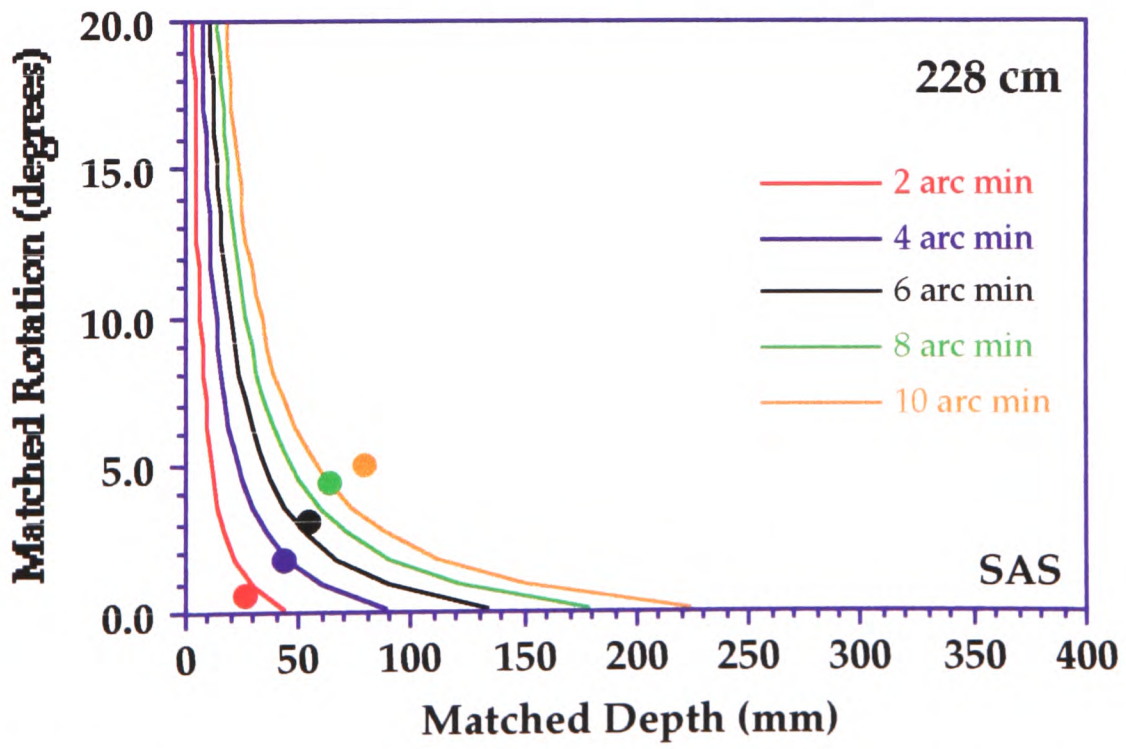


Fig. 5.54: Matched Rotation as a function of relative motion amplitude for proximally identical stimuli of size 10° at two viewing distances, for Subject MML. The dashed lines represent the simulated depth for viewing distances of 114 cm (steeper slope) and 57 cm.

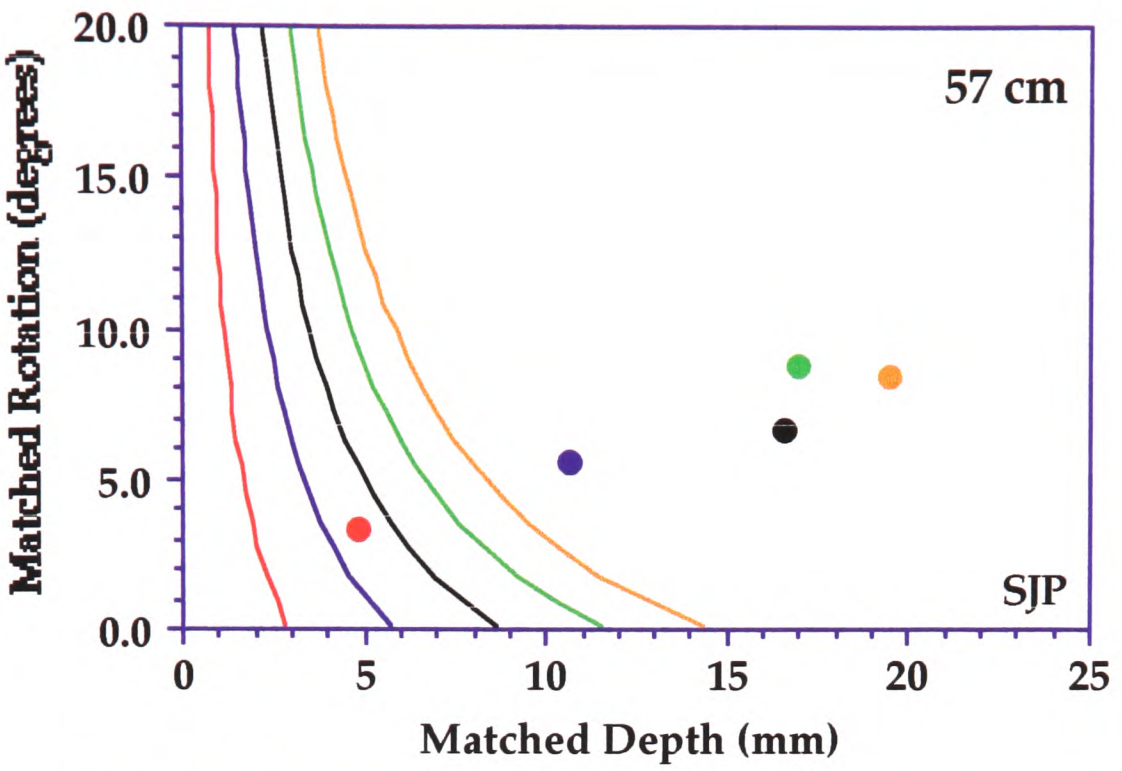
Figs. 5.54 - 5.63 show Matched Rotation against Matched Depth for (proximal) size 5° stimuli of relative motion amplitudes 2 - 10 arc min equivalent disparity. The viewing distance is shown in the top right-hand corner. The depth/rotation combinations predicted by the model are also shown. The legend shows the relative motion amplitude (arc min equivalent disparity) which produces each theoretical depth/rotation curve. Symbols of identical colour arose from surfaces containing identical relative motion amplitudes. The data should theoretically lie on the curve of the same colour. For all surfaces tested, there was an over-estimation of the combined amounts of depth and rotation relative to those predicted by the geometric model. At the nearest viewing distance the perceived depths were over-estimated. At the greatest distance (228 cm), the estimated depth ranged between 50 and 35% of its simulated value, yet the magnitude of perceived rotation was sufficiently great to over-account for the remainder of the relative motion.

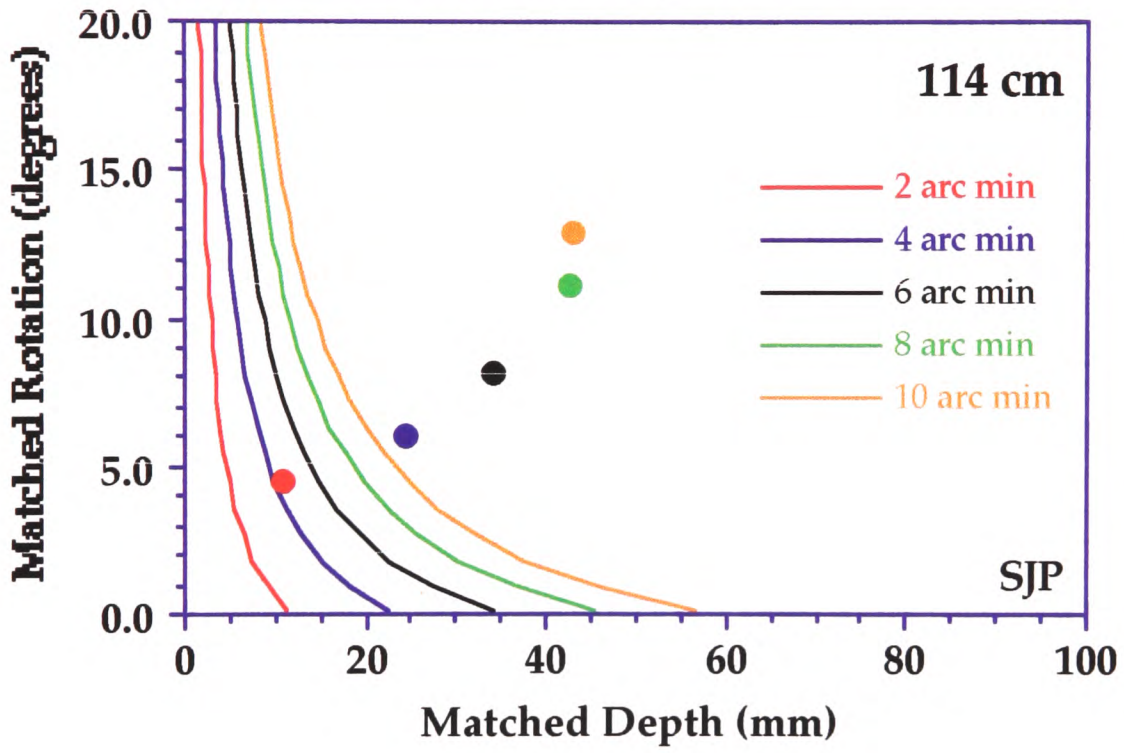


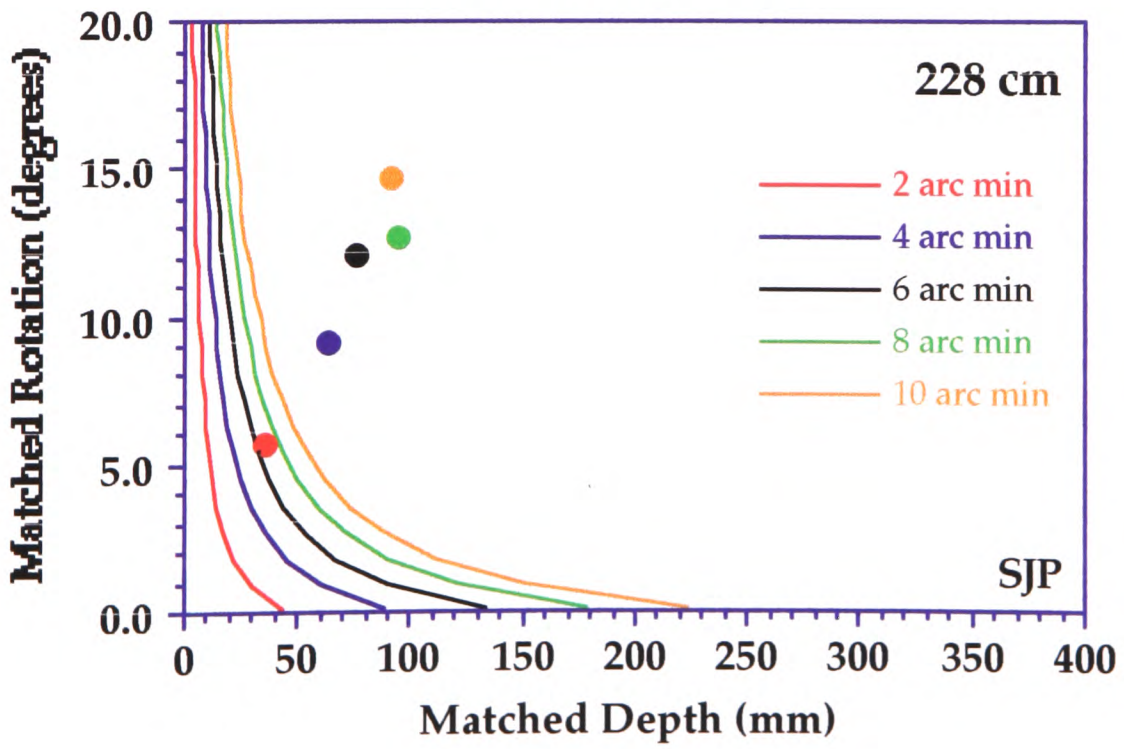




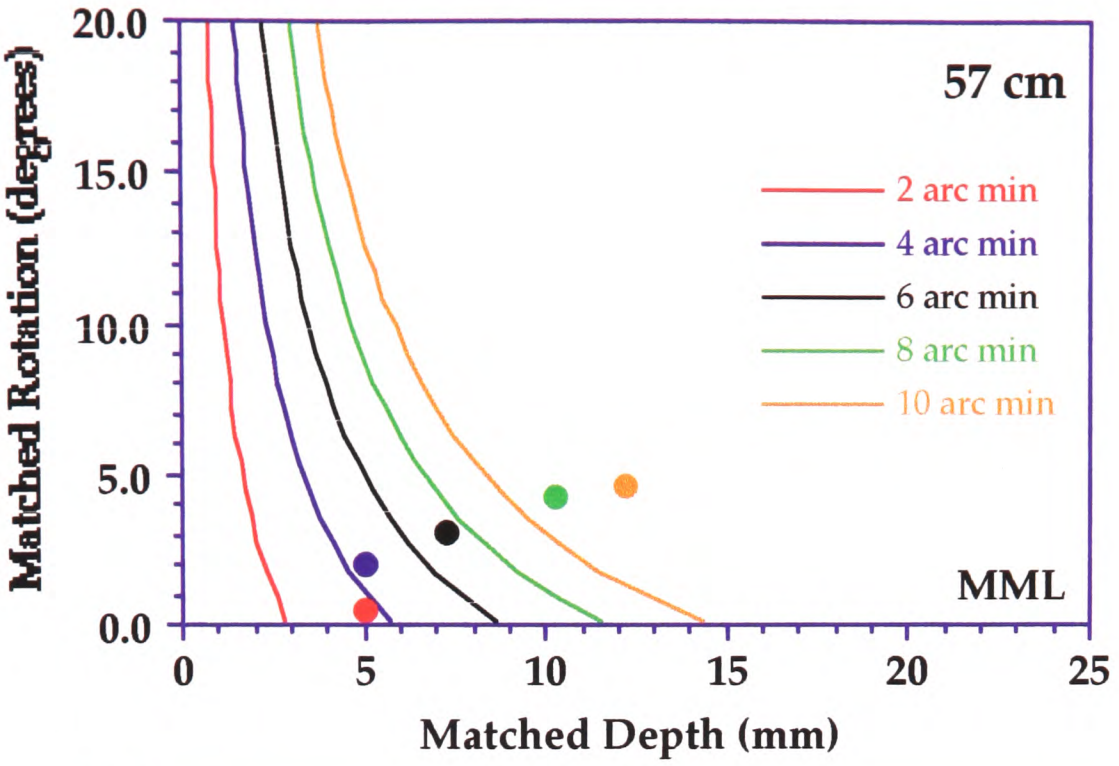
Figs. 5.55 - 5.57: Matched Rotation versus Matched Depth for a 0.4 cpd, size 5° stimulus at three viewing distances for Subject SAS. See text for details.

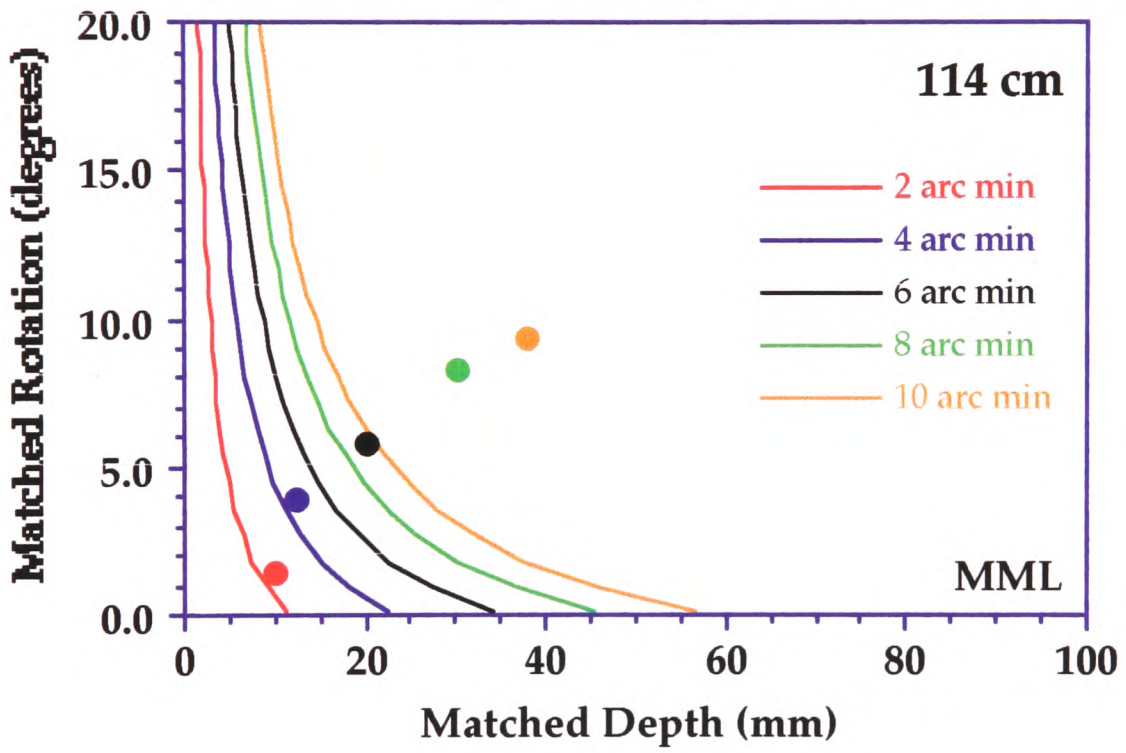


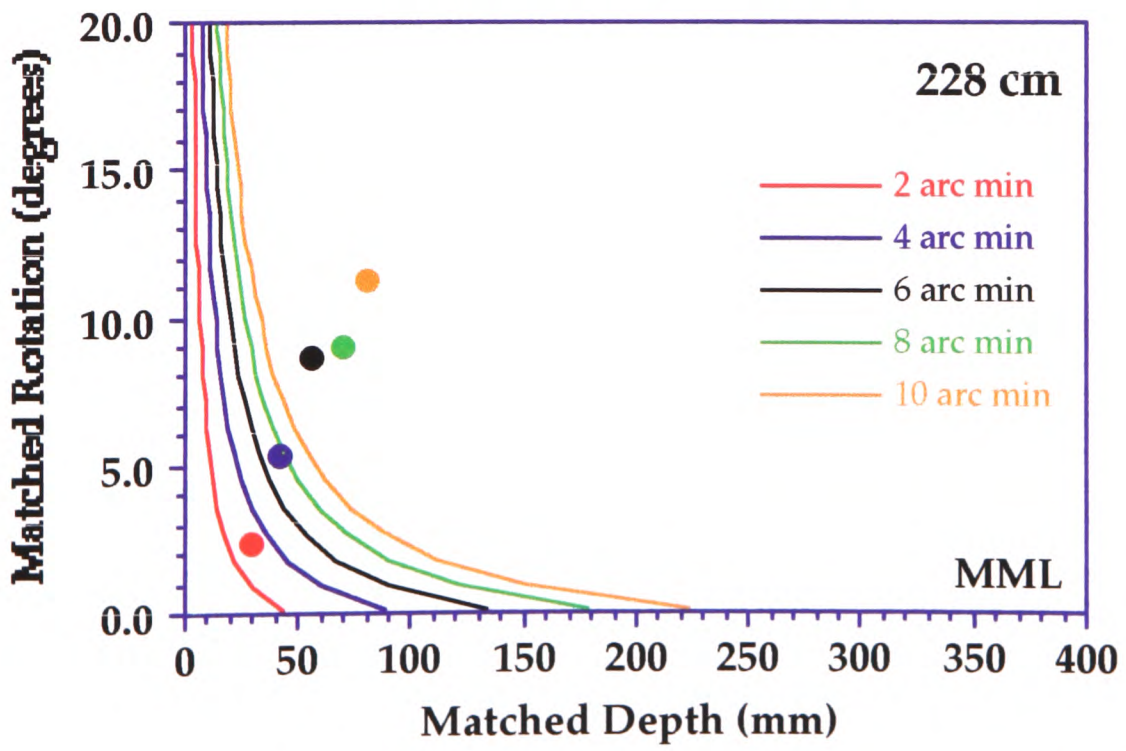




Figs. 5.58 - 5.60: Matched Rotation versus Matched Depth for a 0.4 cpd, size 5° stimulus at three viewing distances for Subject SJP. See text for details.

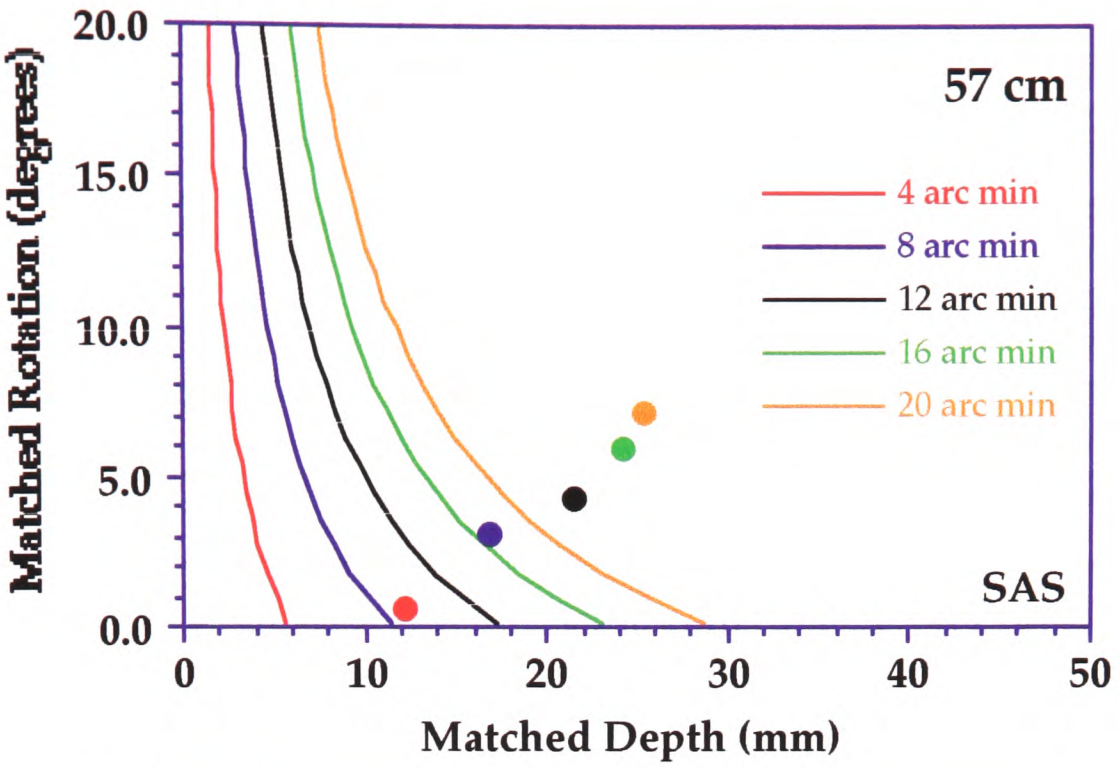


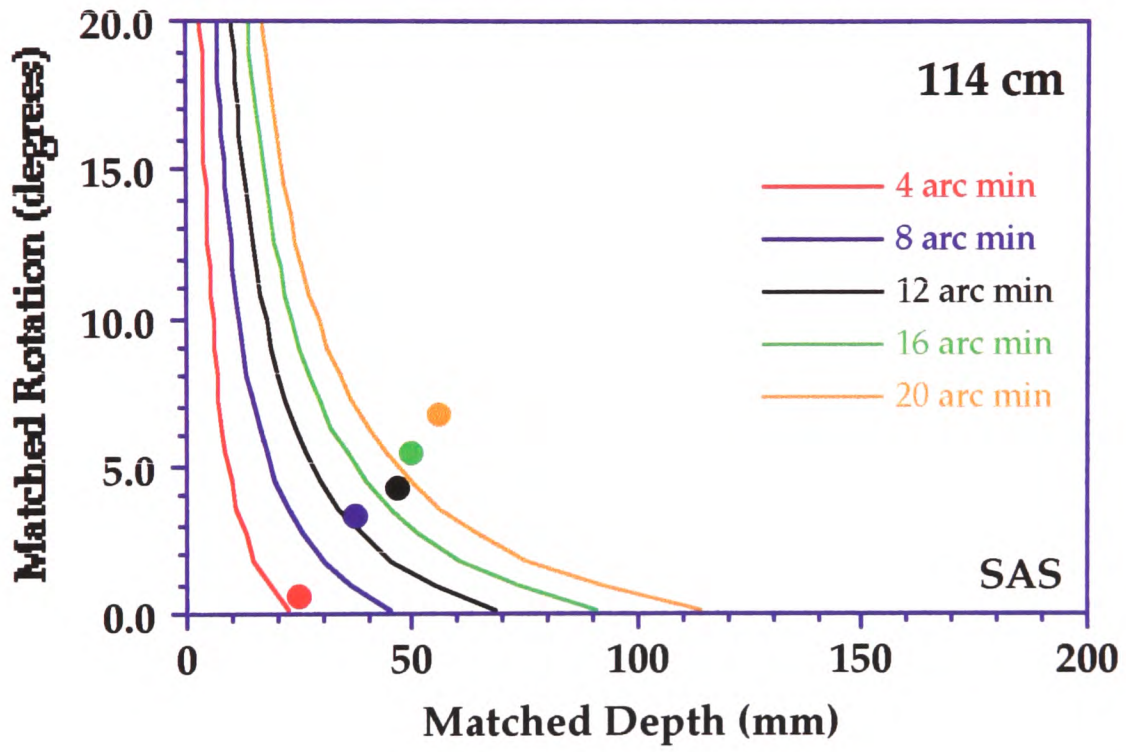




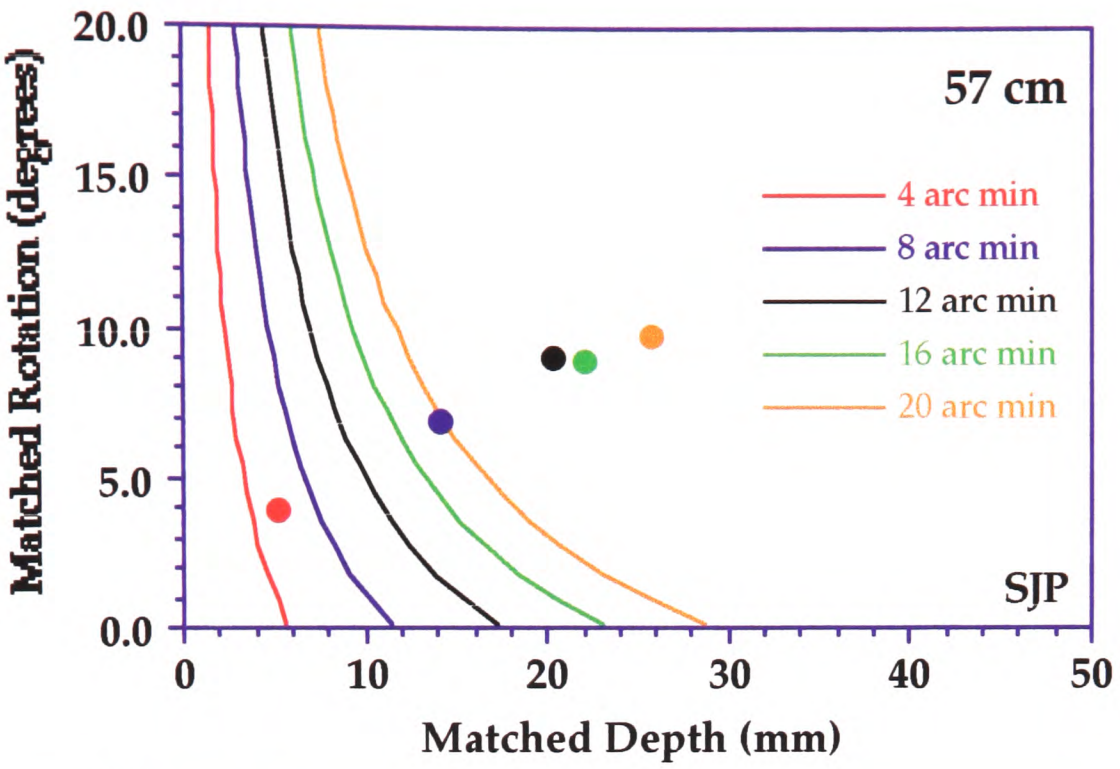
Figs. 5.61 - 5.63: Matched Rotation versus Matched Depth for a 0.4 cpd, size 5° stimulus at three viewing distances for Subject MML. See text for details.

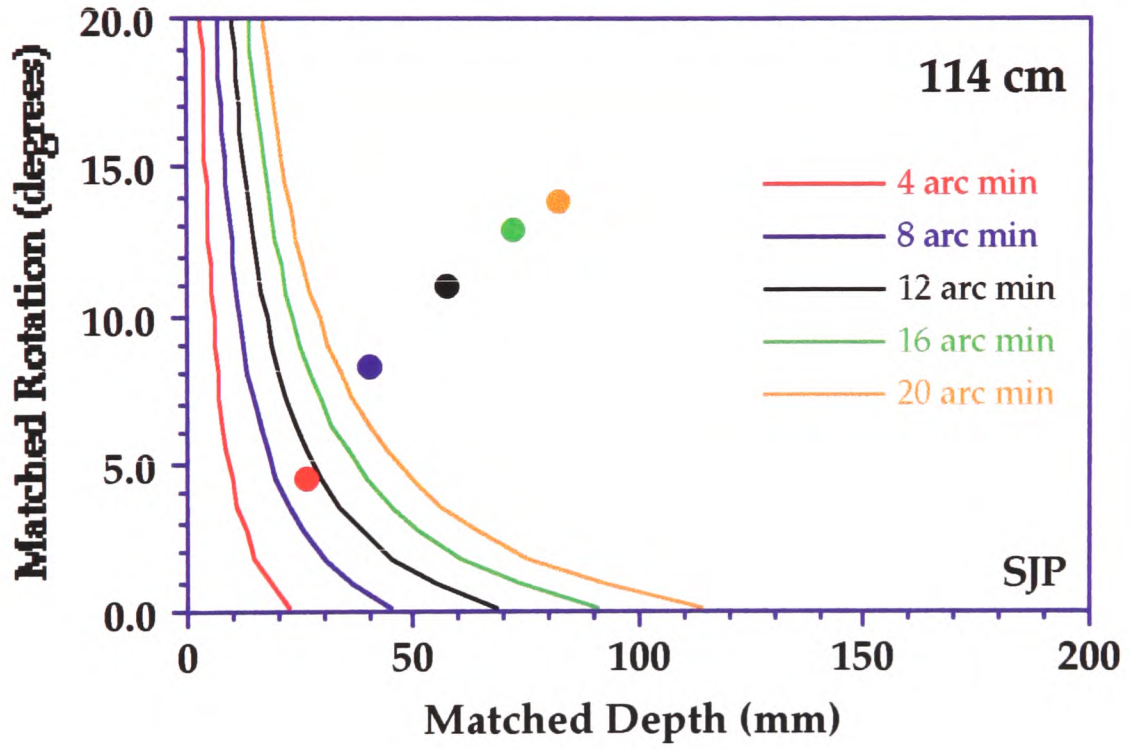
Figs. 5.64 - 5.69 show Matched Rotation against Matched Depth for (proximal) size 10° stimuli of relative motion amplitudes 4 - 20 arc min equivalent disparity. Again, there is an over-estimation of perceived depth at the 57 cm viewing distance. At the distance of 114 cm the depth data lie closer to the ordinate, indicating that the magnitude of perceived depth was less than simulated (assuming zero perceived rotation). Once again, however, the perceived rotation magnitudes were sufficiently large to over-account for the remainder of the relative motion.



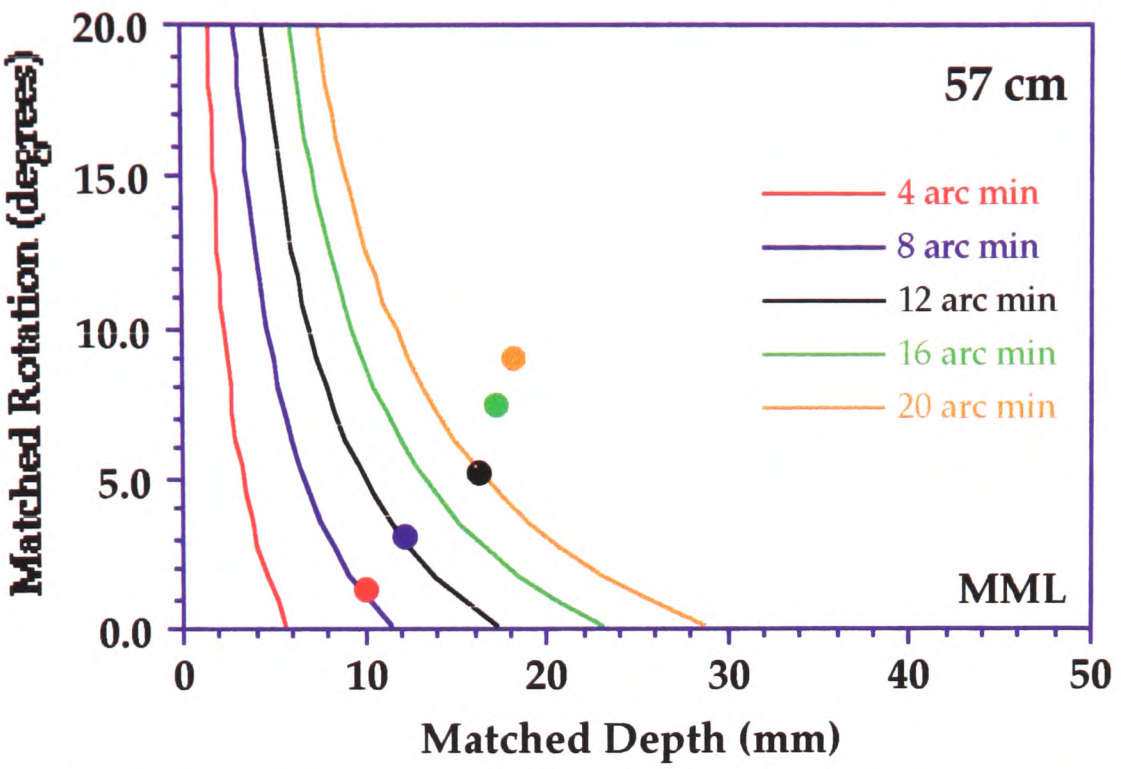


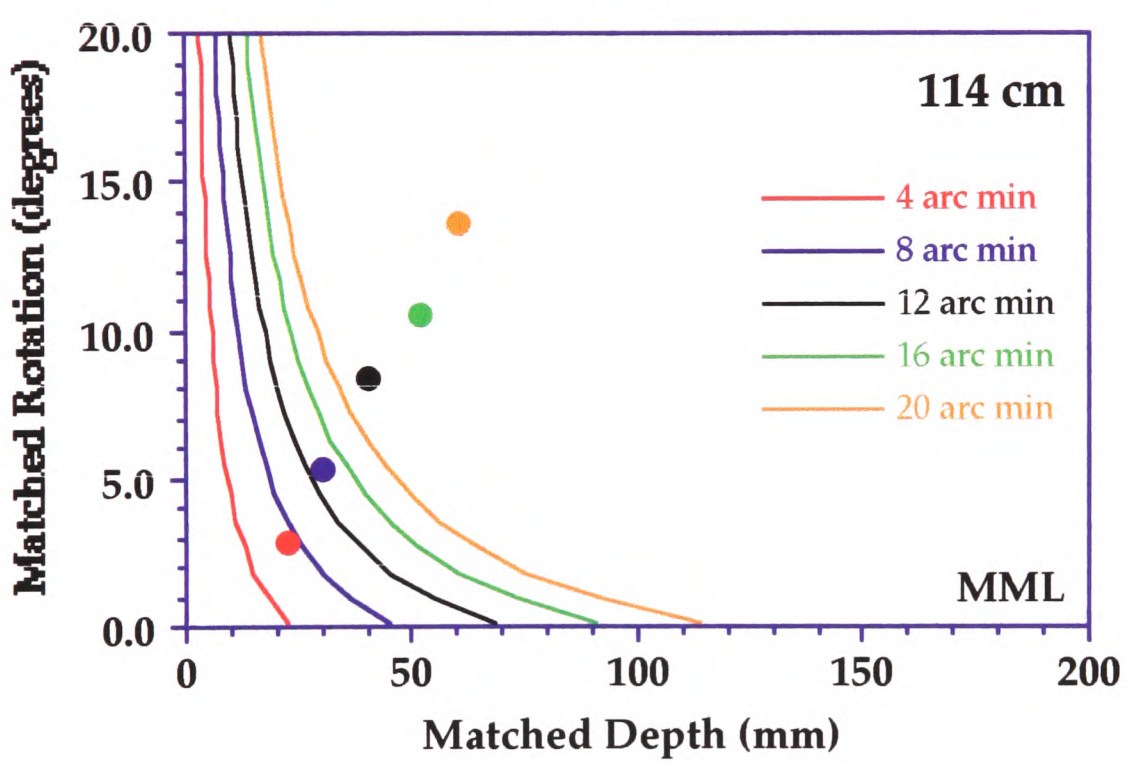
Figs. 5.64 & 5.65: Matched Rotation versus Matched Depth for a 0.4 cpd, size 10° stimulus at two viewing distances for Subject SAS. See text for details.





Figs. 5.66 & 5.67: Matched Rotation versus Matched Depth for a 0.4 cpd, size 10° stimulus at two viewing distances for Subject SJP. See text for details.





Figs. 5.68 & 5.69: Matched Rotation versus Matched Depth for a 0.4 cpd, size 10° stimulus at two viewing distances for Subject MML. See text for details.

5.7.5 Discussion

The finding that perceived rotation increases with increasing viewing distance complements the observation that the Transition Point falls with increasing viewing distance (see Experiment 4). The finding that the perceived depth magnitude also increased with viewing distance is consistent with the scaling of relative motion with absolute distance information. The fact that there is under-estimation at the greater viewing distances concords with other studies which have also found imperfect depth scaling (Ono et al., 1986; Rivest et al., 1989).

Figs. 5.55 - 5.69 show that over-estimation of the depth/rotation combinations predicted by the geometric model occurs at every viewing distance. At the greatest distance (228 cm) the estimated depth ranges between 50 and 35% of its simulated value, yet the magnitude of perceived rotation is sufficiently great to over-account for the remainder of the relative motion.

It is also notable that the relative effects of variation in the relative motion amplitude upon the perceived depth and rotation magnitudes is different at different viewing distances. At 57 cm an increase in the amplitude leads to greater increases in perceived depth (relative to the range of simulated depths under the assumption that the extent of perceived rotation is zero) than in perceived rotation. At 228 cm the converse is true. The bias of the visual system towards the perception of greater amounts of rotation at the further viewing distances may arise as a consequence of the fact that in the real world, the size of the vertical perspective changes accompanying the rotation of a given object through a given angle decreases with increasing viewing distance. The visual system may perhaps have a propensity to rely less heavily on vertical perspective information as a cue to slant as viewing distance increases. In the case of stimuli which remain constant in angular size as viewing distance varies (such as those in the present experiment), the vertical perspective changes accompanying a given rotation angle remain constant also. However, the putative diminished reliance on the vertical perspective information might result in its diminished efficacy at constraining the surface to zero rotation. The effect of vertical perspective on perceived rotation is examined in the following chapter.

An alternative explanation of the data is that the comparatively large depths specified by the parallax at large viewing distances are not perceptually supported by the additional visual information specifying that the display is flat. With increasing distance an increasing amount of the surrounding environment falls within the field of view and the observer may be increasingly affected by such flatness cues. Following the logic of Rogers and Collett (1989), the visual system may perceive the structure whose depth and rotational characteristics conflict least with the information specified by the rest of the visual scene. In the present case this might involve the interpretation of the majority of relative motion as rotation rather than depth.

5.8 Summary and conclusions

The experiments of this chapter examined observers' percepts of depth and rotation, and depth and shearing, in sinusoidal and square wave surfaces respectively. The aim of the experiments was to determine whether the relative magnitudes of these percepts conformed to the predictions of the geometric models of Fig. 5.5 and Fig. 5.8.

Experiments 11 and 12 explored the perceived magnitudes of depth and rotation in sinusoidal surfaces which varied in corrugation frequency (Experiment 11) and size (Experiment 12). It was found that decreases in corrugation frequency and increases in stimulus size each led to increases in the perceived depth magnitude, and to decreases in perceived rotation. These results were interpreted as qualitative support for a trade-off between the extents of depth and rotation. The combined magnitudes of the two attributes were, however, greater than those predicted by the model. Three possible factors were suggested to account for this. These were firstly, an over-estimation of the amount of 2-D relative motion (Ono and Steinbach, 1990; Sato, 1989); secondly, a given extent of relative motion being 'counted' both as depth and subsequently as rotation (or vice versa) and thirdly, possible biases due to the psychophysical technique.

Experiments 13 and 14 investigated perceived depth and shearing in square wave surfaces. Experiment 13 employed surfaces of 0.05 and 0.4 cpd, and found that a greater extent of depth was perceived in the surface of

lower corrugation frequency. The magnitude of perceived shearing was greater for the 0.4 cpd surface. These data are once again suggestive of a trade-off relationship between the extents of perceived depth and perceived motion of a surface. Again, the combined magnitudes of depth and shearing were greater than those predicted by the geometric model, although the extent of the over-estimation was less for these surfaces than for the sinusoidal surfaces of Experiment 11. The extent of perceived shearing was invariably greater than the physical 2-D relative motion amplitude, supporting the notion that at least a portion of the over-estimations observed in Experiments 11 and 12 might be attributable to an over-estimation of the extent of 2-D relative motion. Experiment 14 used stimulus sizes of 2.5° and 20°. The perceived extents of both depth and shearing were greater for the larger surface, a finding obviously inconsistent with the predictions of the geometric model. At higher relative motion amplitudes there was a decrease in the extent of perceived depth for both stimulus sizes — however, the combined extents of perceived depth and shearing were again greater than predicted.

Experiment 15 examined perceived depth and rotation in proximally identical stimuli at viewing distances of 57, 114 and 228 cm in order to determine whether the lack of depth constancy in parallax surfaces at distances greater than about 1 m (Ono et al., 1986; Rivest et al., 1989) coexists with increased magnitudes of perceived rotation. This was found to be the case. Despite the fact that depth was increasingly under-estimated with increasing viewing distance, the increased extents of perceived rotation were sufficient to over-account for the amount of relative motion.

Chapter 6:

The Effect of Dynamic Vertical Perspective on Perceived Rotation and Perceived Depth in Motion Parallax Surfaces

6.1 Introduction

The word 'perspective' refers to the differences in the geometric properties of the projected retinal images which arise as a consequence of variations in viewing distance (Howard and Rogers, 1995). Perspective information is commonly sub-divided into two categories: horizontal perspective and vertical perspective. Fig. 6.1 below provides a simple instance of each type of perspective information.

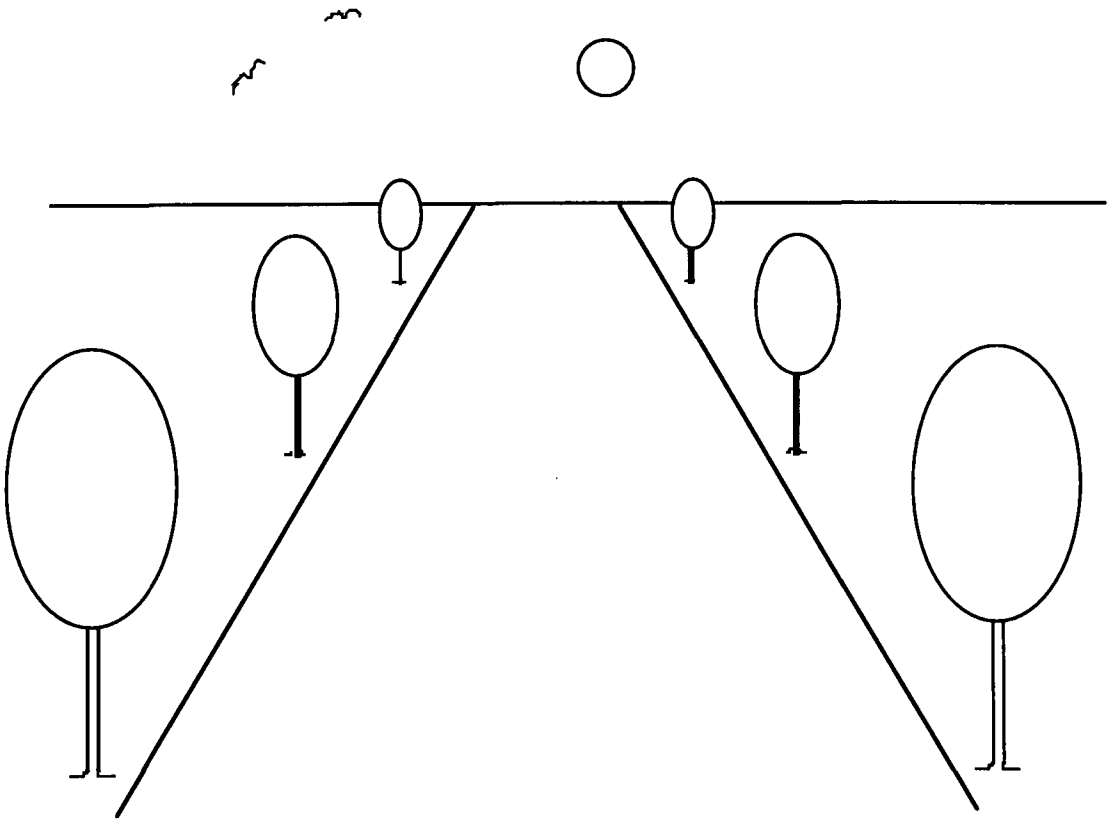


Fig. 6.1: The decrease in the projected horizontal size of the road is interpreted as arising from its increased distance from the observer. This is an instance of horizontal perspective information. The decrease in the projected vertical size of the trees is interpreted as arising from their increased distance from the observer. This is an instance of vertical perspective information.

Horizontal perspective refers to the differences in the horizontal angular subtense of an object's projected image due to variations in the distance separating object and observer. Thus in the diagram above the decrease in the projected horizontal size of the road, for instance, is interpreted as arising from its increased distance from the observer. Vertical perspective refers to the difference in the vertical angular subtense of an object's projected image due to variations in the distance separating object and observer. In the diagram above, the decrease in the projected vertical size of the trees is interpreted as arising from their increasing distance from the observer.

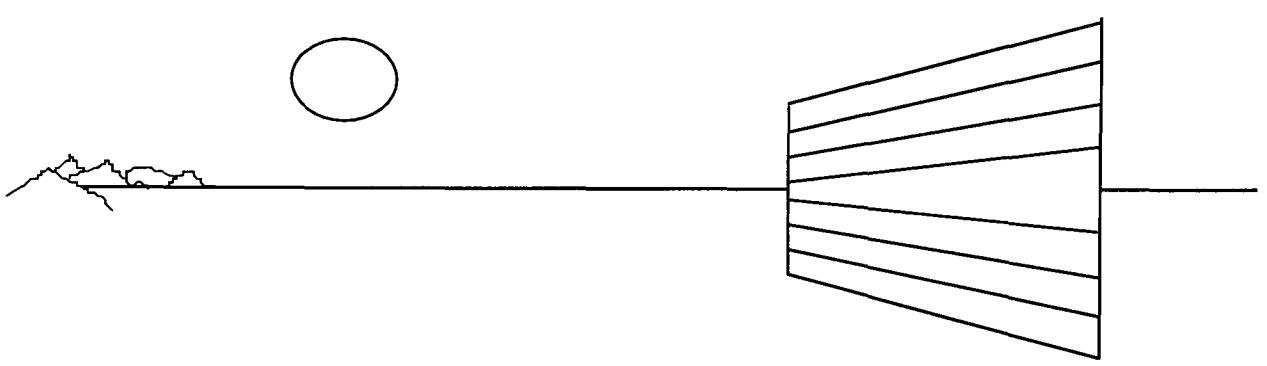


Fig. 6.2: The vertical edge subtending the greater vertical extent is interpreted as being nearer to the observer than the vertical edge subtending the smaller vertical extent. The surface is thus perceived to be slanted (with respect to the observer) about a vertical axis .

As a consequence of exactly the same principle, the angles subtended by the vertical edges of a surface which is slanted, with respect to the observer, about a vertical axis will differ. The nearer edge of the surface will subtend a larger angle than the further edge, given that the physical sizes of the edges are equal (Fig. 6.2).

When discussing perspective information it is important to specify the direction of slant of the surface in question. For instance, the term 'foreshortening' is used to describe both the shortening of the vertical lines on a surface inclined about a horizontal axis, and the shortening of the horizontal lines on a surface slanted about a vertical axis as the surface inclination or surface slant increases. The term 'texture gradient' refers to the decreasing size and increasing 'bunching' of the images on an inclined or slanted surface which occur with increasing distance. On a curved surface, variations in depth also cause variations in the aspect ratio of texture elements. For instance, the images of circles on the external surface of a convex cylinder become increasingly elliptical at the curve of the cylinder. Cutting and Millard (1984) and Todd and Akerstrom (1987) have

shown that the most effective information about the recession of a planar surface is provided by a simple texture gradient, whereas the most useful source of information about surface curvature is provided by the aspect ratio.

Both Figs. 6.1 and 6.2 provide striking examples of linear perspective, the convergence of physically parallel lines with increasing distance. When the term is used to refer to the convergence of the horizontal contours of a slanted surface (Fig. 6.2), it is an example of vertical perspective information — it is the horizontal spatial gradient in vertical image size. The term can also be used to describe the *vertical* spatial gradient in *horizontal* image size — horizontal perspective information. In Fig. 6.1, use of the term 'linear perspective' is equally valid in the description of both the convergence of the sides of the road, and the convergence of the imaginary contours joining the tops of the trees, as each recede into the distance. The former is an instance of horizontal perspective, the latter an instance of vertical perspective. For this reason, as noted above, it is important to define the direction of surface slant. This thesis uses surfaces which are slanted about the vertical axis, and thus the latter instance of linear perspective is more relevant. Since the term is potentially confusing, the term 'vertical perspective' rather than 'linear perspective' will generally be used. Where the term 'linear perspective' is used in the following discussion of the literature, it refers to the effects of vertical, rather than horizontal, perspective.

Vertical perspective is an effective cue to the slant of both rotating surfaces (Gibson and Gibson, 1957; Braunstein and Payne, 1968; Börjesson, 1971; Hershberger, Stewart and Laughlin, 1976; Braunstein, 1977; Rogers and Rogers, 1992) and stationary surfaces (Epstein and Mountford, 1963; Clark, Smith and Rabe, 1966; Youngs, 1976; Stevens and Brookes, 1988; Gillam and Ryan, 1992; Eby and Braunstein, 1993; Ryan and Gillam, 1994) surfaces. The influence of perspective information on the perceived *direction* of rotation (and thus perceived depth order) is considered in Section 6.8.1. The current discussion deals with the perceived *magnitude* of rotation or slant.

Gibson and Gibson (1957) presented observers with continuous perspective projections of a surface which rotated about a vertical axis

away from the frontal plane to a slant angle of 15 - 70°. Observers were asked to indicate the magnitude of perceived slant using the adjustable pointer on a protractor. Estimates were found to be accurate, both for regular and irregular shapes. When the same judgement was made using stationary stimuli which were merely presented at a given slant angle, performance was much impaired, especially for the irregular stimuli. Whilst not specifically addressing the question of the relative importance of horizontal and vertical perspective, this work demonstrates that the dynamic perspective cue can result in veridical slant perception regardless of the type of pattern carrying the transformation. Williams (1993) showed that the perceived rotation magnitude of a surface defined by motion parallax and binocular disparity was strongly influenced by the rotation angle specified by a dynamic vertical perspective cue, and furthermore that for surfaces in which stereo and vertical perspective cues specified the opposite depth order, the latter cue could determine the perceived order, provided that the disparity cue was not too large.

Clark et al., (1955) showed observers outlines of real rectangular film forms which were slanted at angles of either 0, 20 or 40° and were viewed monocularly through a reduction screen. These stimuli gave rise to perceived slant which, though underestimated, led the authors to conclude that linear perspective is a sufficient cue to slant in the absence of other visual cues.

Gillam (1968) used aniseikonic lenses to introduce binocular disparities into surfaces containing perspective slant cues, reasoning that "... the more effective the perspective cues present, the less should be the susceptibility of a surface to aniseikonic distortion". In a preliminary experiment, subjects monocularly viewed slanted surfaces which contained one of four types of information: foreshortening, linear perspective, the two types of information or irregular gradients (the stimulus markings consisted of randomly scattered dots). The slant angle was 6 - 24°. Observers adjusted a Meccano wheel to the apparent slant of the surface. Slant perception was most accurate in the presence of linear perspective, and was least accurate in conditions in which the surface contained foreshortening alone or irregular gradients. The addition of foreshortening to linear perspective did not improve performance. When the surfaces were viewed in the frontal plane with the aniseikonic lens, significantly less aniseikonic slant

was perceived in surfaces which contained effective perspective cues than in those which did not, demonstrating that linear perspective could successfully compete with the stereo information, and that perceived slant magnitude was based on a compromise between the two sources of information with neither completely dominating the percept. More recently, using large (>60 degrees) fields, Gillam (1998) has shown that strong perspective information can result in stereoscopic slant reversals in situations in which both the horizontal and vertical image sizes are magnified in one eye.

Youngs (1976) investigated the importance of stereoscopic and perspective information in slant perception, and additionally whether the perceived slant of surfaces (either rectangles or trapezoids) was different from the perceived slant of stimuli which contained solely the vertical contours of the surfaces. Firstly, it was found that stimuli containing perspective information were judged to be significantly more slanted than those which did not. The addition of disparity to stimuli containing perspective did not result in a significant difference in the magnitude of perceived slant. Secondly, there was no significant difference between the magnitudes of slant perceived in the surface stimuli and in the contour stimuli, suggesting that perspective information can be as effectively utilised in the latter as in the former.

Eby and Braunstein (1993) examined the roles of compression and linear perspective in the perception of real, slanted surfaces. It was found that under conditions of no cue conflict, linear perspective was the most effective cue to slant. The slants of surfaces consisting of compression alone were underestimated by about 40%.

The above evidence demonstrates that vertical perspective is a particularly important determinant of the perceived slant of both stationary surfaces and translating, or rotating, surfaces.

Binocular vertical perspective cues can also provide information concerning the shape and depth of stereoscopic surfaces. Despite the fact that the majority of the disparities between the two eyes' images are in the horizontal dimension, small differences also exist between the vertical positions of the image components. These arise because features which are

not in the median plane are at different distances from each eye and thus project to different vertical positions. Vertical disparities can theoretically provide an estimate of viewing distance (Longuet-Higgins, 1982; Mayhew and Longuet-Higgins, 1982; Gillam and Lawergren, 1983). Note that there are alternative cues to viewing distance — for example oculomotor cues (Foley, 1980; Cormack, 1984) and pictorial cues such as familiar size or perspective (O' Leary and Wallach, 1980; Predebon, 1993). If distance is known, then the horizontal disparities can be depth-scaled.

Initial investigations found that the manipulation of vertical disparities did not affect the perceived depth or local shape of stereoscopic surfaces (Sobel and Collett, 1991; Cumming, Johnston and Parker, 1991). However, subsequent studies found that vertical disparity manipulations can affect the perceived size, depth and shape of stereoscopic surfaces (Rogers and Bradshaw, 1992, 1993, 1995; Bradshaw et al., 1996; Lipson, Rogers and Ledgeway, 1998). It is likely that the reason for the earlier studies' failure to find an effect of vertical disparity is the relatively small display sizes used — Cumming et al.'s (1991) display, for instance, subtended only about 11°. Bradshaw and Rogers (1994) investigated depth scaling as a function of display size, and showed that scaling was zero when the display was 10° in diameter. Additionally, vergence and accommodation were held constant in the study of Sobel and Collett (1991), and thus provided absolute distance cues which conflicted with the vertical disparity information. Bradshaw, Glennerster and Rogers (1996) examined the effect on disparity scaling of vergence manipulations (keeping vertical perspective constant) and vertical perspective manipulations (keeping vergence constant). The influence of vergence was greatest in the smallest display, whereas vertical perspective was effective only in the large (>20°) displays. For each display size the effect of the two cues was approximately additive. This study confirms that both oculomotor and vertical perspective cues are used to scale horizontal disparities, and illustrates that the relative importance of these cues is dependent upon the size of the display.

In the motion parallax domain, a frontal surface located in the median plane will contain a changing spatial gradient of vertical size as the observer makes side-to-side head movements. As noted by Howard and Rogers (1995), when the observer moves through the interocular distance,

the oppositely-signed spatial gradients of vertical size which exist at the beginning and at the end of the observer's translation are analogous to those which exist simultaneously in binocular viewing. The authors additionally note that, since the change in spatial gradient which occurs with the observer's translation specifies the change in obliqueness of the surface with respect to the eye, if the distance through which the eye has translated is known then the absolute distance to the surface can be calculated. Absolute distance information can be calculated in the binocular viewing case if the separation of the eyes is known.

There is no evidence suggesting that the changes in the vertical size gradient which accompany head movements are used to estimate the absolute distance of surfaces. Rogers and Bradshaw (1991), however, reported that changes in the amounts of perceived depth and rotation of parallax surfaces occur in response to vertical perspective manipulations.

The following experiments examine the influence of dynamic vertical perspective on the magnitudes of perceived rotation and depth in head-linked motion parallax surfaces. Experiments 11 and 12 showed that when the perceived rotation magnitude of a surface was increased, either through increasing the corrugation frequency or decreasing the stimulus size, the magnitude of perceived depth was decreased. These results are qualitatively in line with those predicted by the geometric model, which illustrates the causal relationship between the relative magnitudes of these variables. At first sight the results of Experiments 11 and 12 seem to suggest that an increase in perceived rotation would *necessarily* lead to a decrease in perceived depth, regardless of the factor causing the variation in perceived rotation magnitude. However, the fact that the trade-off relationship has been found when rotation has been manipulated by variations in corrugation frequency or stimulus size does *not* necessarily mean that it exists when alternative cues are used to manipulate perceived rotation. The following experiments attempt to determine whether the trade-off relationship occurs when the latter is manipulated via a dynamic vertical perspective cue.

6.2 Experiment 16: The Effect of Dynamic Vertical Perspective on Matched Rotation and Matched Depth as a function of Corrugation Frequency

6.2.1 Apparatus

In order to simulate the vertical perspective changes created by a surface rotating about a vertical axis, the display underwent trapezoidal transformations in synchrony with the lateral head movements of the observer. As the head reached the endpoint of its translation, the size of the nearer vertical side of the display was maximal and the size of the opposite vertical side was minimal. As the head moved to the other endpoint, the transformation occurred in the opposite direction. The transformation was achieved by modulating the size of the signal sent to the Y deflection plates of the scope according to a ramp waveform synchronised to the line rate of the display. The waveform was generated by a Wavetek 186. Further details of the apparatus and the display calibration are found in Section 2.1.4.

Note that the simulation of the rotation was an approximation to the perspective distortion created by a rotating surface, and that horizontal perspective effects such as width changes and texture density gradients did not accompany the vertical perspective manipulations (see Section 2.1.4).

6.2.2 Stimuli

Stimuli had corrugation frequencies of 0.05, 0.1, 0.2 or 0.4 cpd, and were of size 20°. Relative motion amplitude was 8, 16, 24, 32 or 40 arc min equivalent disparity. There were two vertical perspective levels. The simulated rotations with respect to the screen were 0° and ±15° for a ±6.5 cm head translation. These corresponded to rotations with respect to the eye of ±6.5° and ±21.5° for the ±6.5 cm head movement (Fig. 6.3). The direction of rotation was opposite to the direction of head movement.

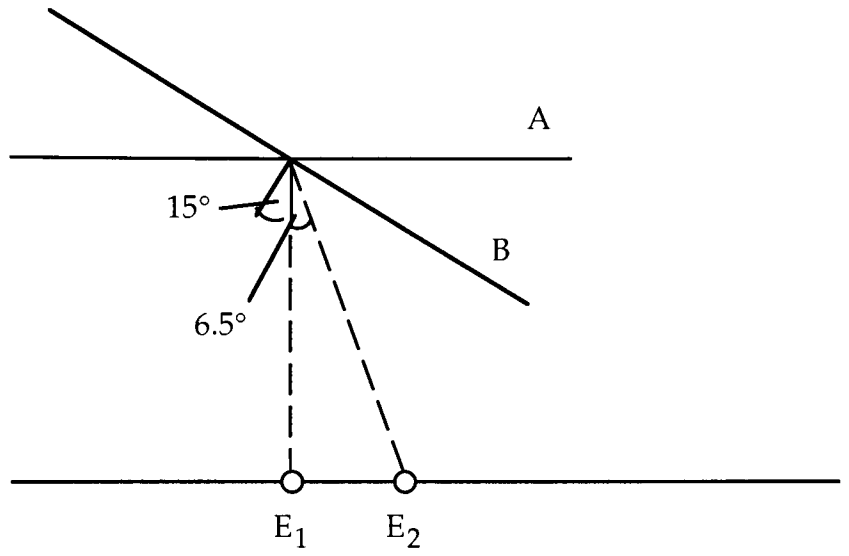


Fig. 6.3: As the eye moves laterally through 6.5 cm from position E_1 to position E_2, the surface rotates from position A, through 15°, to position B. The rotation of the surface with respect to the eye is 6.5° + 15° = 21.5°.

6.2.3 Procedure

Each trial required the subject to make a Depth Match and a Rotation Match. Trials were blocked in groups of ten, corresponding to two repetitions of each of the relative motion amplitudes. Subjects completed three sessions, each consisting of eight blocks (four corrugation frequencies x two levels of vertical perspective). The order of blocks within each session was randomised. In half the total number of trials the depth judgement was made first and the rotation judgement was made second. In the other half, this order was reversed.

Matched Rotation was measured in the same way as in the previous experiments of this thesis. The technique used to obtain Matched Depth values differed from that used in Experiments 11 - 14. Pilot experiments showed that the depth matching data obtained with the electronic callipers under conditions of vertical perspective changes were not clearly separable from those obtained under conditions of no vertical perspective changes. This was true for surfaces of corrugation frequency 0.05, 0.1 0.2 and 0.4 cpd. The data are shown in Appendix D.

A more sensitive technique of measuring perceived depth was introduced. This technique employed an electronic match stimulus. This stimulus was identical to the test stimulus except that it did not contain the dynamic vertical perspective cue: only horizontal relative motion was present. The degree of this parallax motion was controlled by a potentiometer, operated by the observer and situated out of sight. The test

and match stimuli alternated every 2 seconds with an inter-stimulus interval (ISI) of approximately 0.2 seconds, during which time the relative motion amplitude was zero. The onset of the test stimulus was marked by a tone. The onset of the match stimulus was marked by a tone of lower frequency. Observers were instructed to alter the relative motion amplitude of the match stimulus until the depth of the two stimuli appeared identical. Test and match stimuli alternated until the observer was satisfied with the match.

In half of the trials the rotation match was made first, followed by the depth match. In the other half this order was reversed. Six rotation matches and six depth matches were made for each of the eight stimulus conditions.

6.2.4 Results

Figs. 6.4 - 6.7 show Matched Rotation (left-hand panels) and Matched Depth (right-hand panels) as a function of relative motion, for each corrugation frequency for Subjects SAS, SJP and MML. The open symbols represent matches obtained when the magnitude of rotation simulated by the vertical perspective cue was 0°. The closed symbols represent matches obtained when the magnitude of rotation simulated by the vertical perspective cue was 15°.

The magnitude of perceived rotation increased with increasing corrugation frequency and with increasing relative motion amplitude, as demonstrated in Experiments 7 and 11. For each corrugation frequency, a greater amount of rotation was perceived when the vertical perspective cue signalled a rotation of 15° than when it signalled a rotation of 0°. The range of perceived rotations in the former condition generally spanned 15°. The Matched Rotation functions were approximately parallel under the two vertical perspective conditions. The elevation of the 15° function relative to the 0° function was approximately 10° on average.

For the two lowest corrugation frequencies, Matched Depth was similar in both vertical perspective conditions. At the two higher frequencies, 0.2 and 0.4 cpd, the magnitude of perceived depth was slightly lower when the vertical perspective cue indicated a rotation of 15° than when it indicated a

rotation of 0°.

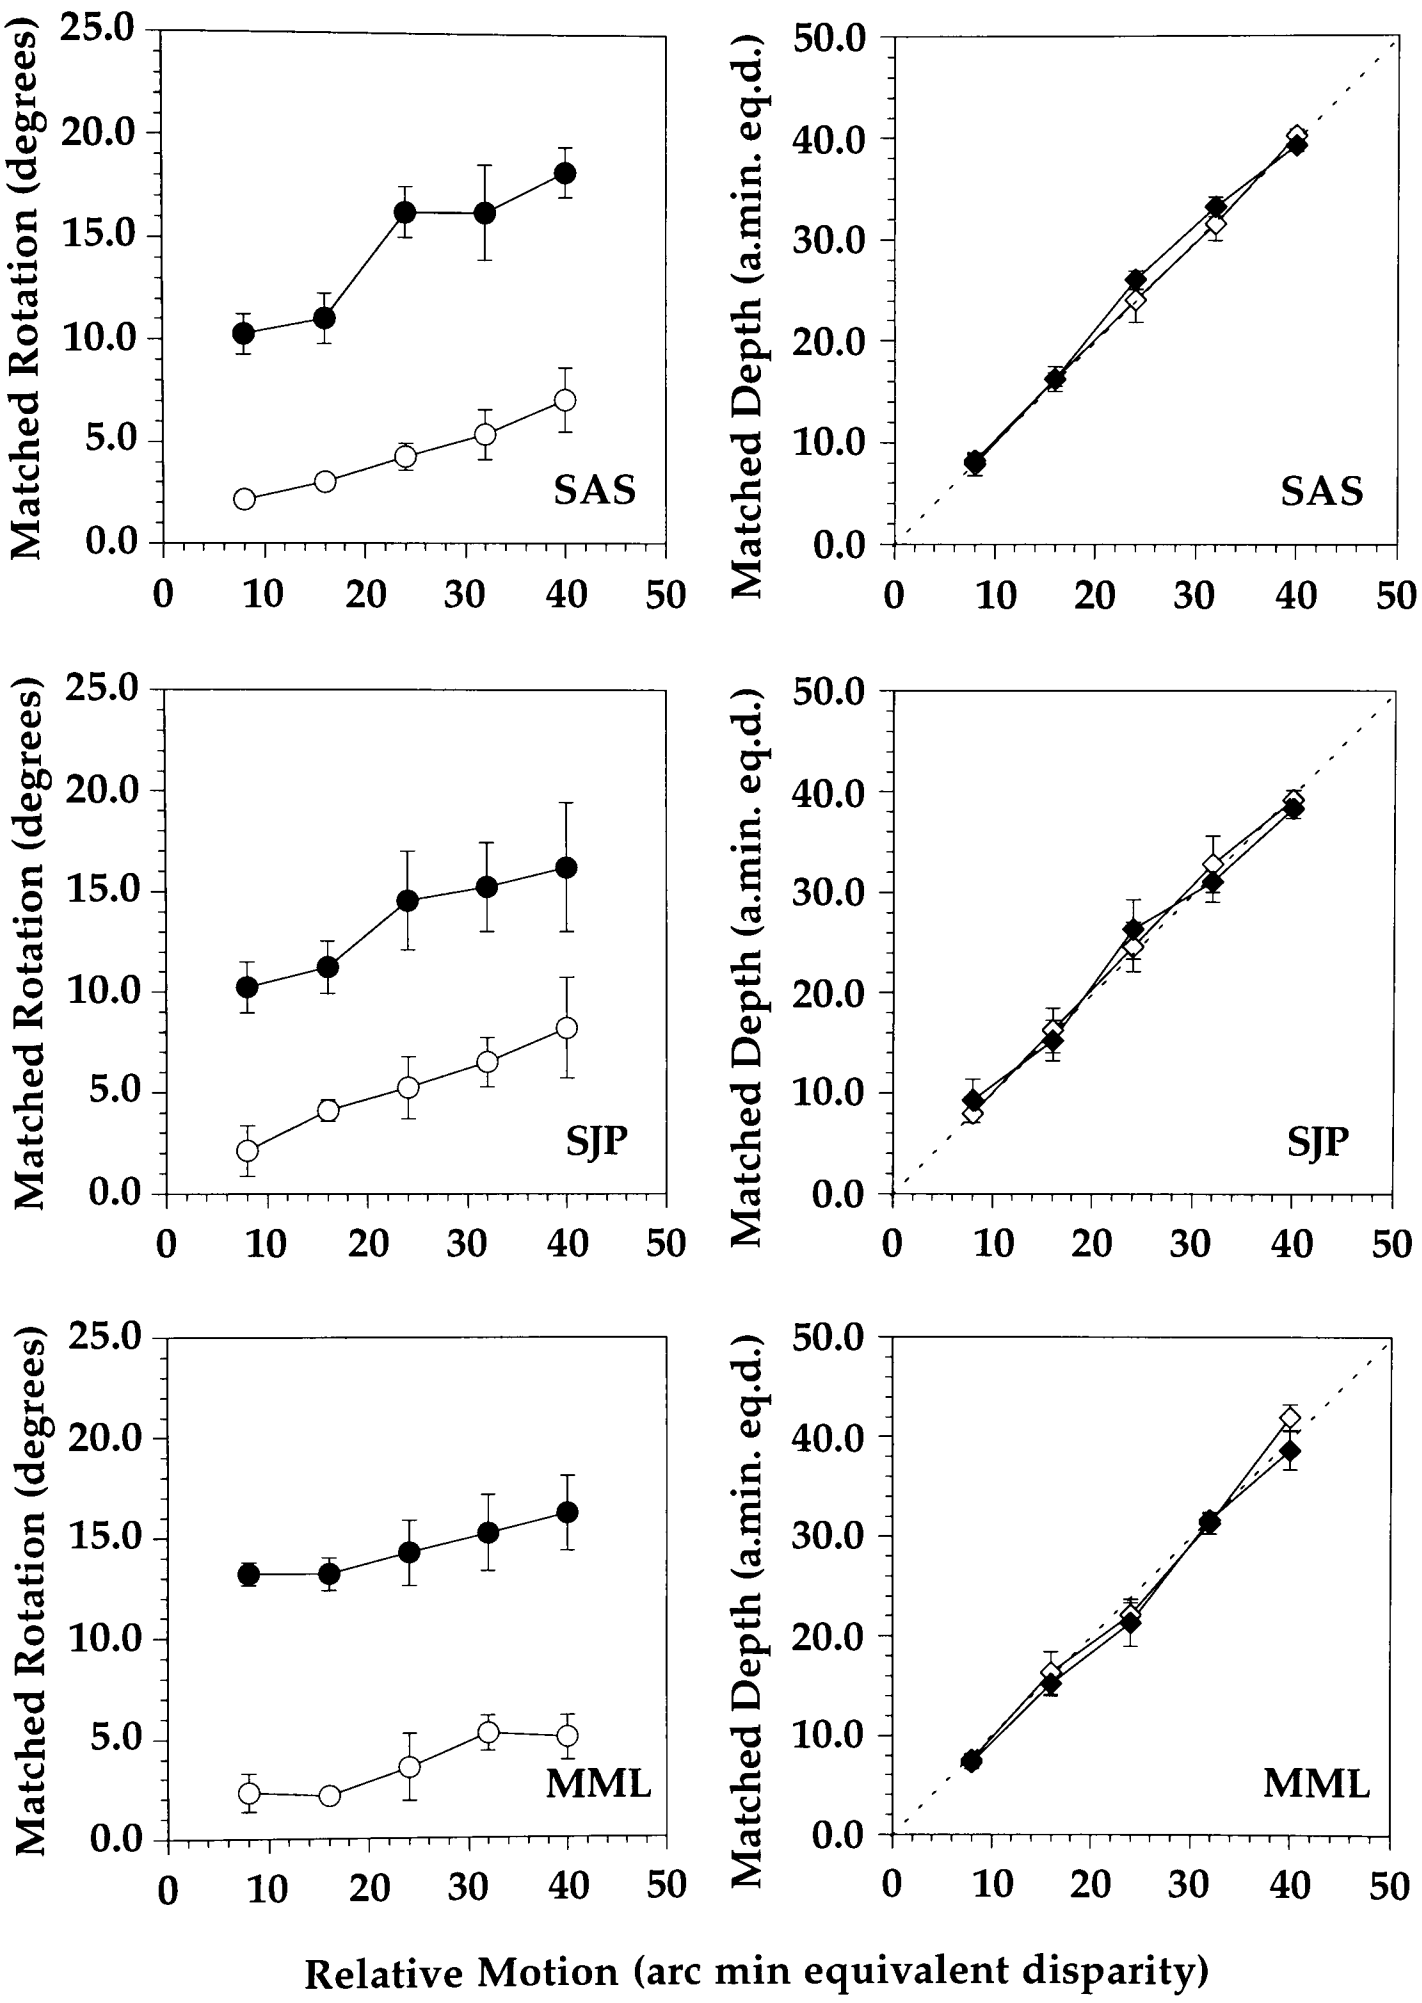


Fig. 6.4: Matched Rotation and Matched Depth for a 0.05 cpd surface of size 20° for Subjects SAS, SJP and MML. The open symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 15°.

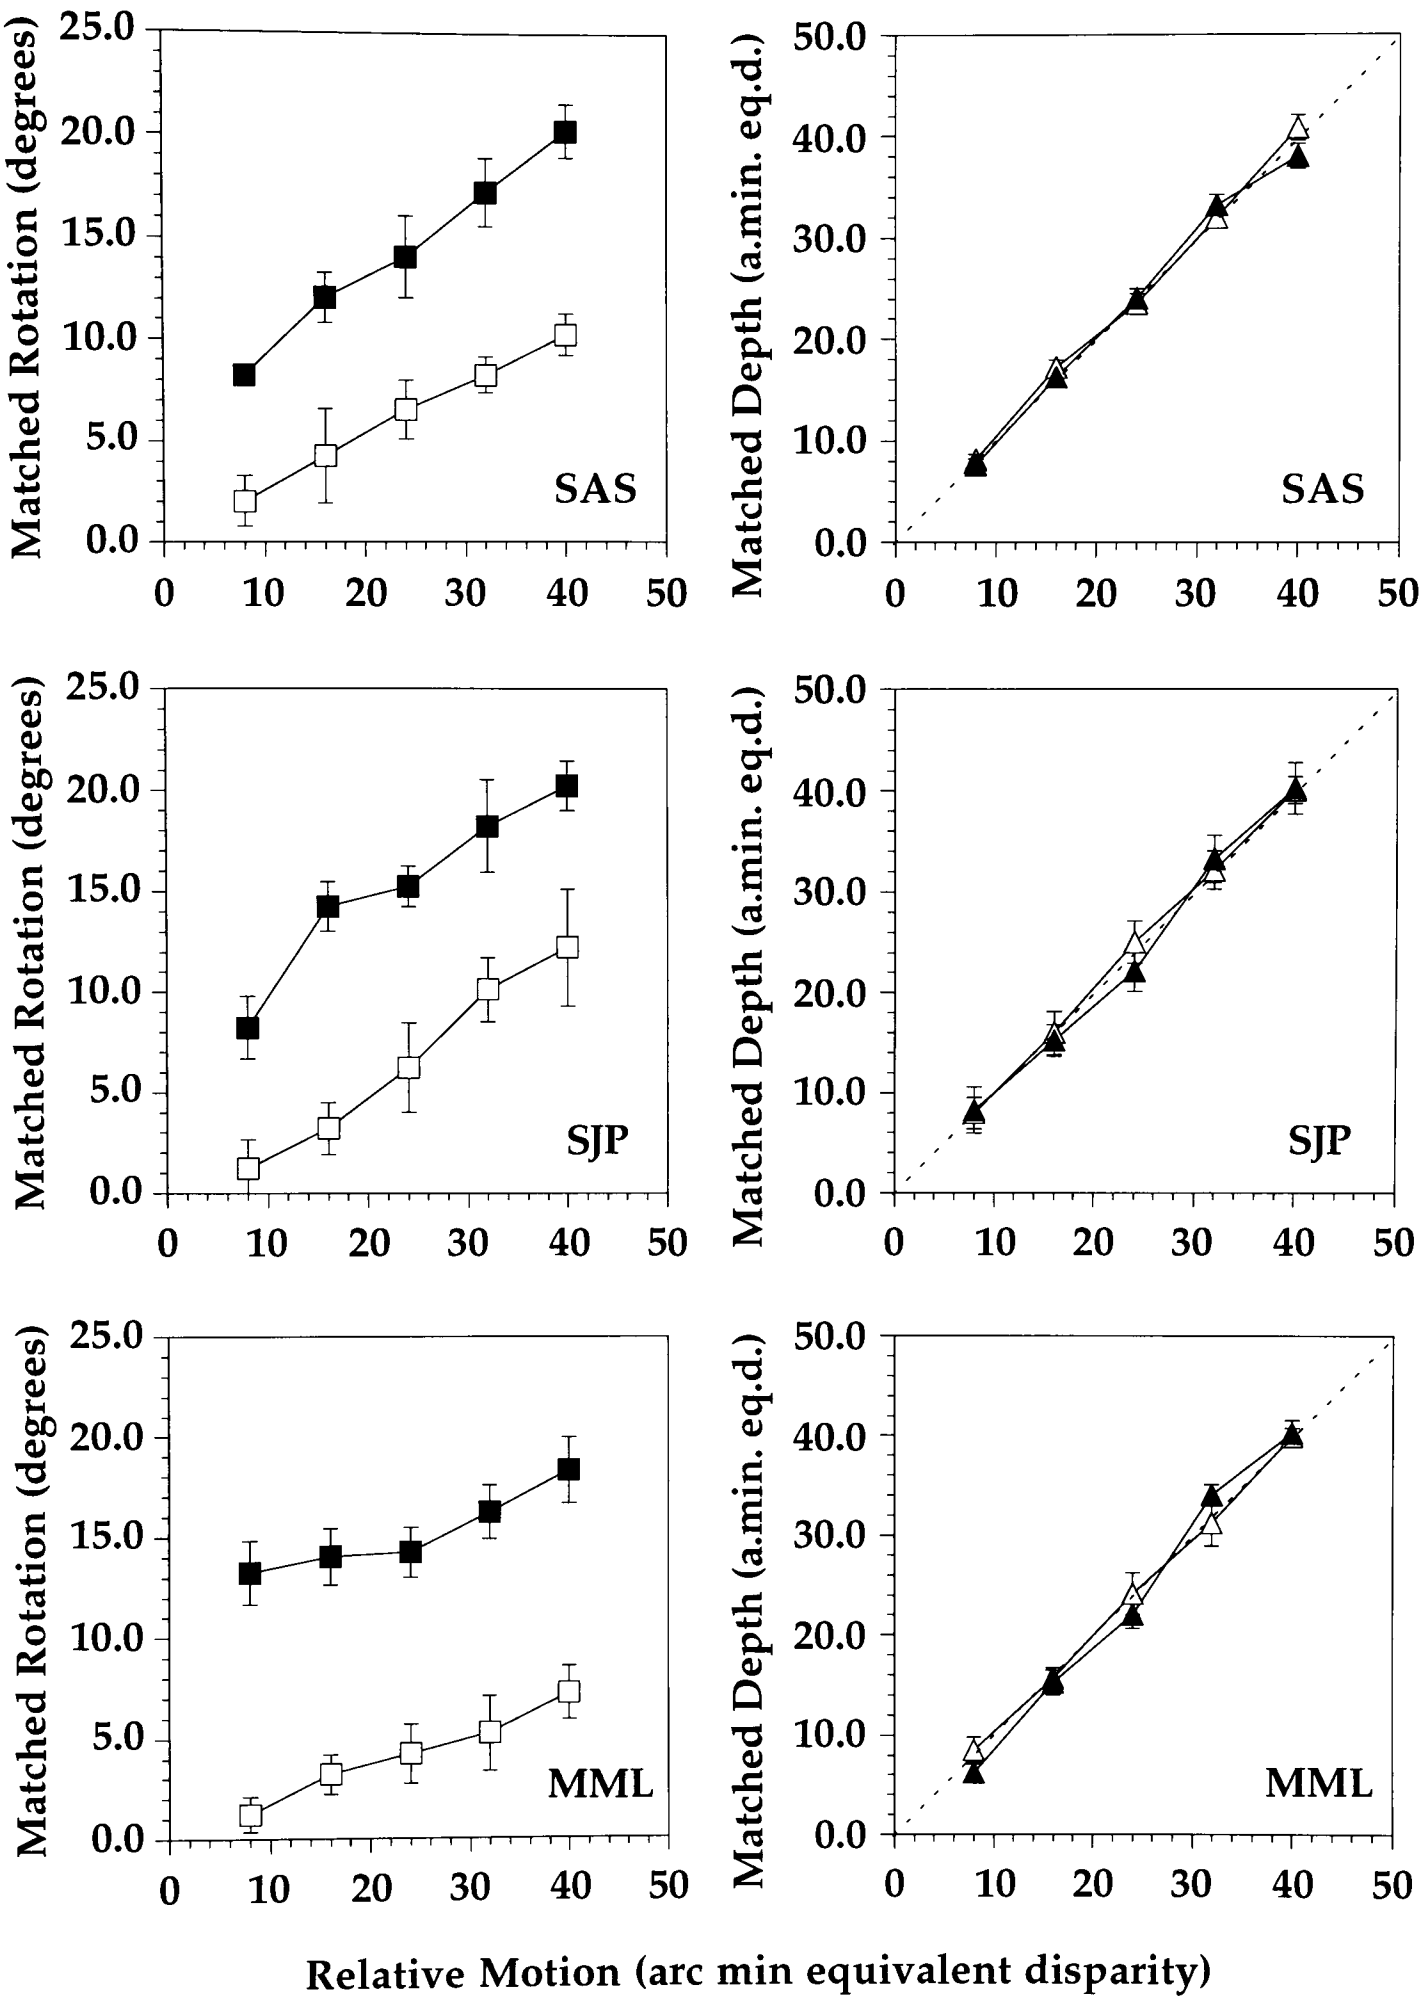


Fig. 6.5: Matched Rotation and Matched Depth for a 0.1 cpd surface of size 20° for Subjects SAS, SJP and MML. The open symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 15°.

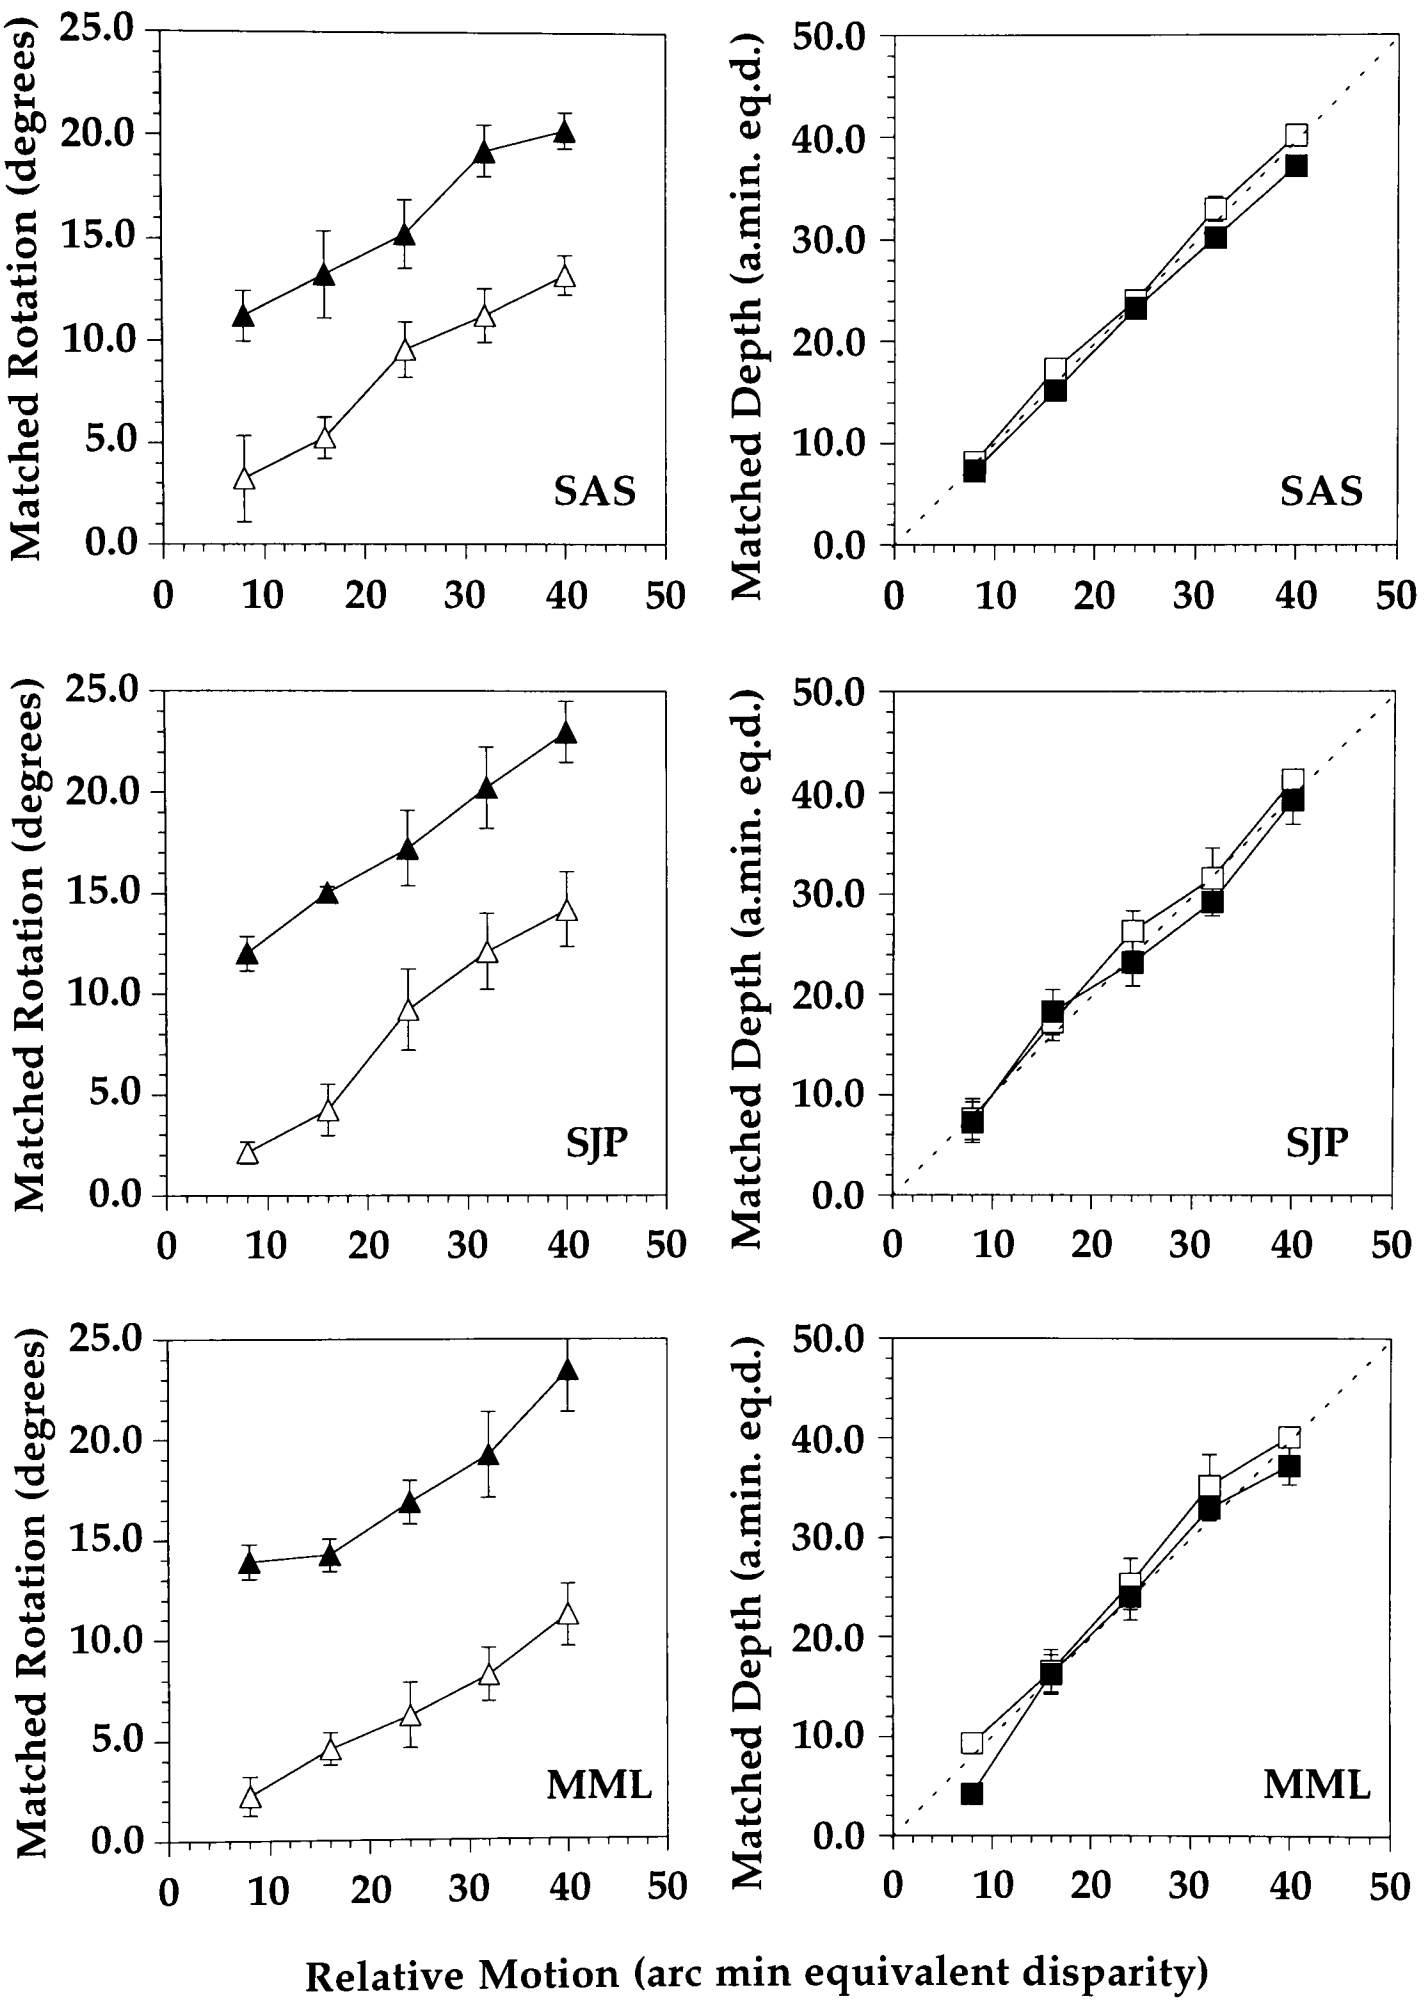


Fig. 6.6: Matched Rotation and Matched Depth for a 0.2 cpd surface of size 20° for Subjects SAS, SJP and MML. The open symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 15°.

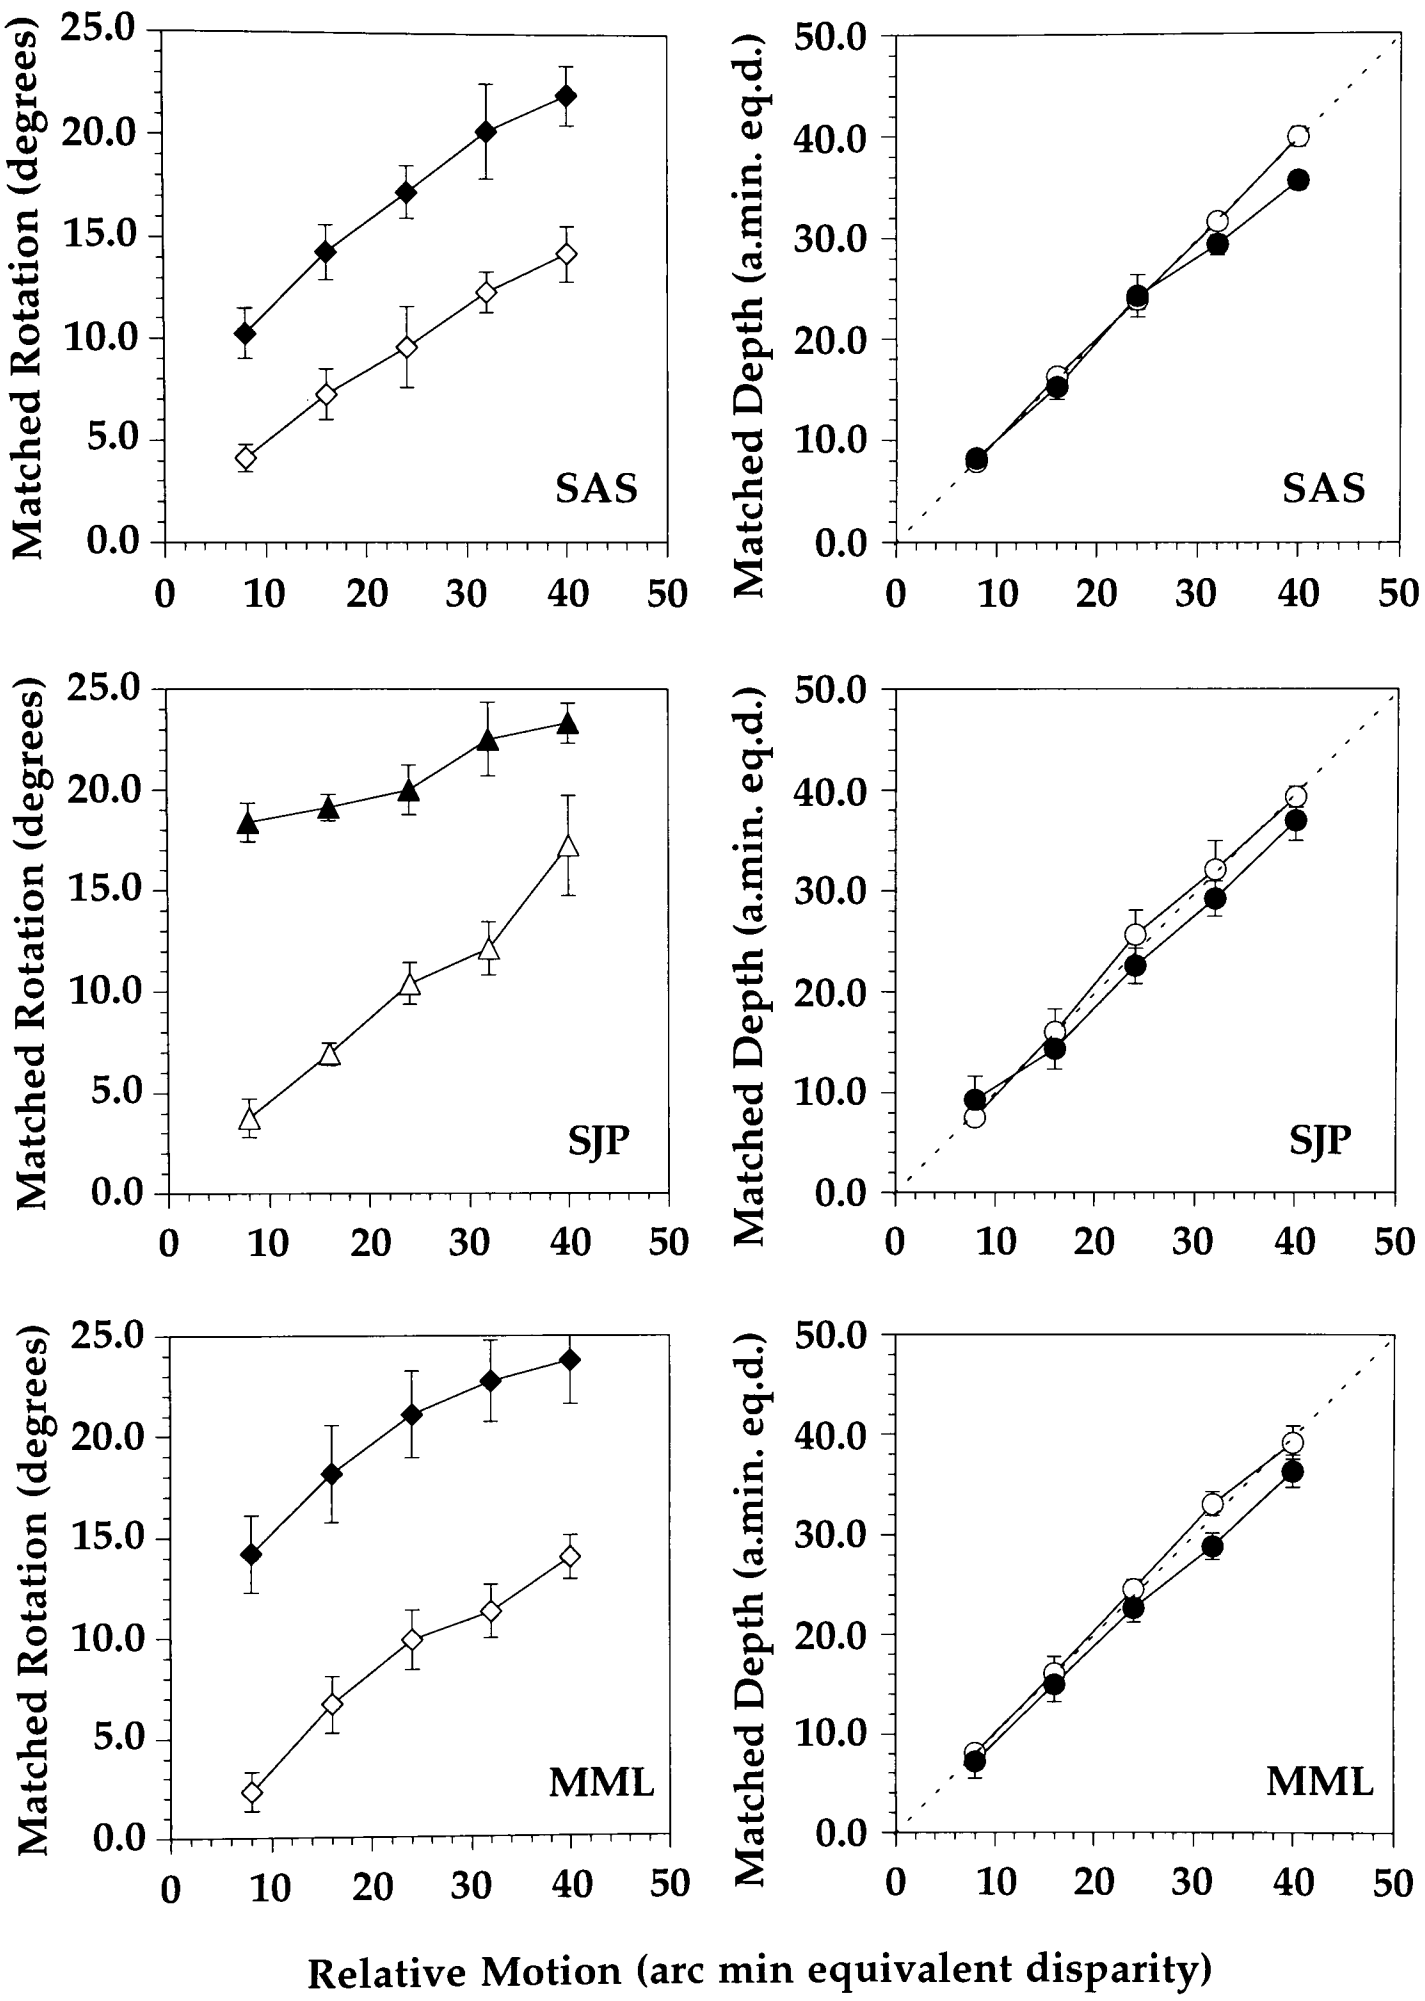


Fig. 6.7: Matched Rotation and Matched Depth for a 0.4 cpd surface of size 20° for Subjects SAS, SJP and MML. The open symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 15°.

6.2.5 Discussion

Vertical perspective was shown to be a strong cue to rotation as predicted, with a consistently greater amount of rotation perceived in the 15° than in the 0° condition. Despite the considerable elevation of perceived rotation at the higher vertical perspective level, the cue does not completely determine the extent of perceived rotation, as indicated by the fact that the Matched Rotation functions are not horizontal. When the cue simulated a rotation angle of 15°, perceived rotation increased with increasing relative motion amplitude. This indicates that the relative strengths of the vertical perspective and relative motion cues in the determination of perceived rotation remain approximately constant when the relative sizes of the cues are varied.

The finding that the large variations in the perceived rotation magnitude were not accompanied by large variations in perceived depth stands in contrast to the results of Experiments 11 and 12. Those experiments demonstrated sizeable trade-offs between perceived rotation and depth. The variations in perceived rotation were achieved through the manipulation of corrugation frequency (Experiment 11) and stimulus size (Experiment 12). The current data indicate that variations in perceived rotation caused by the manipulation of the vertical perspective cue do not result in concomitant variations in depth. This implies that the perception of increased rotation *per se* does not necessarily result in a decreased amount of perceived depth.

It is noted that the depth matching data of the current experiment were collected using a different matching technique from Experiments 11 and 12. It could thus be argued that legitimate comparisons between the two sets of data may not be made. However, it will be recalled that pilot data for the current experiment were collected using electronic callipers and that differences between the depth matches made in the presence and absence of vertical perspective changes were not evident using this technique. Thus the basic finding of the present experiment — namely that vertical perspective changes, while substantially altering the magnitude of perceived rotation, are not accompanied by changes in the magnitude of perceived depth — would *not* appear to be due to the different depth matching technique employed.

Rogers and Bradshaw (1991) reported that manipulations of vertical perspective affected the perceived depth in monocular motion parallax displays. These authors used an electronic depth match stimulus similar to that of the current experiment, and found that the amount of perceived depth in a stimulus in which the vertical perspective cues indicated a rotation of ±22° (with respect to the eye) was approximately 3 arc min equivalent disparity less than the amount of perceived depth in a stimulus in which the vertical perspective cues indicated a rotation of 0°. The size and corrugation frequency of the surfaces of Rogers and Bradshaw (1991) were 21° and 0.2 cpd respectively. The magnitude of the trade-off is quantitatively comparable to that of the similar surface (0.2 cpd, 20°) in the present experiment. It is evident that in both cases the trade-off is small.

The suggestion that increases in perceived rotation magnitude are not necessarily accompanied by decreases in perceived depth is discussed further in Section 6.4. The results of the current experiment suggest that the depth/rotation trade-off is partly dependent upon the nature of the rotation cue. The following experiment examined the effect of vertical perspective on the trade-off as a function of stimulus size.

6.3 Experiment 17: The Effect of Dynamic Vertical Perspective on Matched Rotation and Matched Depth as a function of Stimulus Size

6.3.1 Stimuli

Stimulus size was 5, 10 or 20°. The vertical perspective cue conditions were identical to those above. The corrugation frequency of the stimuli was 0.4 cpd.

6.3.2 Procedure

This was identical to the procedure of the previous experiment.

6.3.3 Results

Figs. 6.8 - 6.10 show Matched Rotation and Matched Depth as functions of relative motion, for the two levels of vertical perspective. Each set of graphs shows results for a single subject.

Each of the graphs in Fig. 6.11 shows Matched Rotation as a function of relative motion for the three stimulus sizes, in order that the relative magnitudes of perceived rotation may be compared. The left-hand column shows matches made in the condition in which vertical perspective simulated a rotation angle of 0°. The right-hand column shows matches made in the condition in which vertical perspective simulated a rotation angle of 15°.

Figs. 6.8 - 6.10 illustrate that the magnitude of perceived rotation was higher when the vertical perspective cue simulated an angle of 15° than when it simulated an angle of 0°. This is consistent with the results of the previous experiment. At the largest stimulus size, 20°, the Matched Rotation functions were generally steeper when the vertical perspective cue signalled 0° rotation than when it signalled a 15° rotation. At the smallest stimulus size, 5°, the functions were approximately parallel. Subject SAS generally tended to perceive less rotation than Subjects SJP and MML, especially in the surfaces for which the simulated rotation angle was 0°.

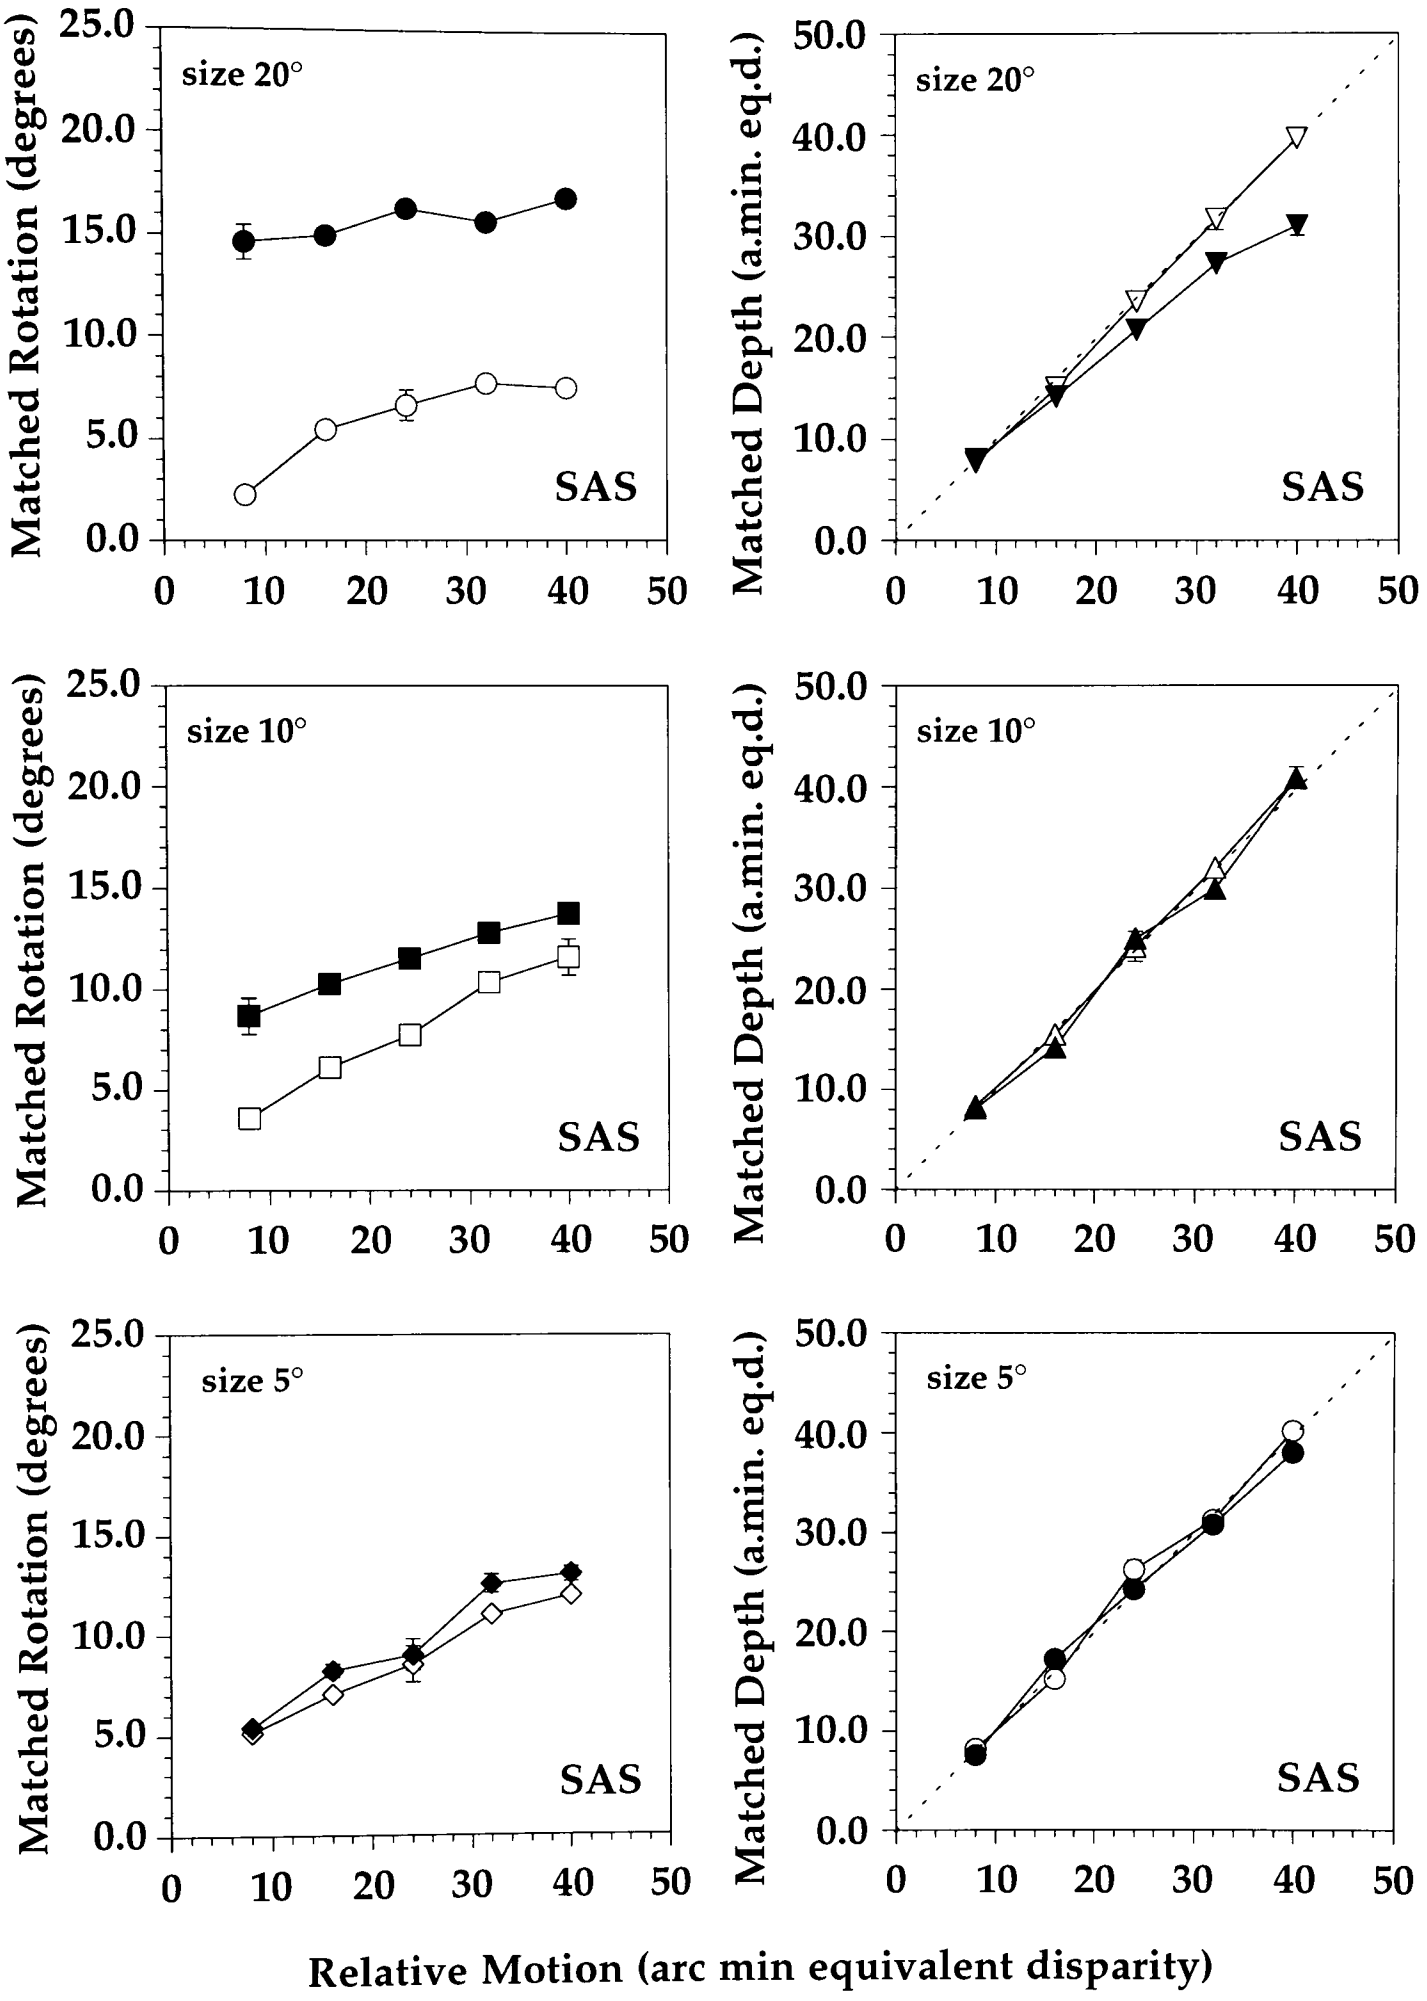


Fig. 6.8: Matched Rotation (left-hand panel) and Matched Depth (right-hand panel) for 0.4 cpd surfaces of size 5°, 10° and 20° for Subject SAS. The open symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 15°.

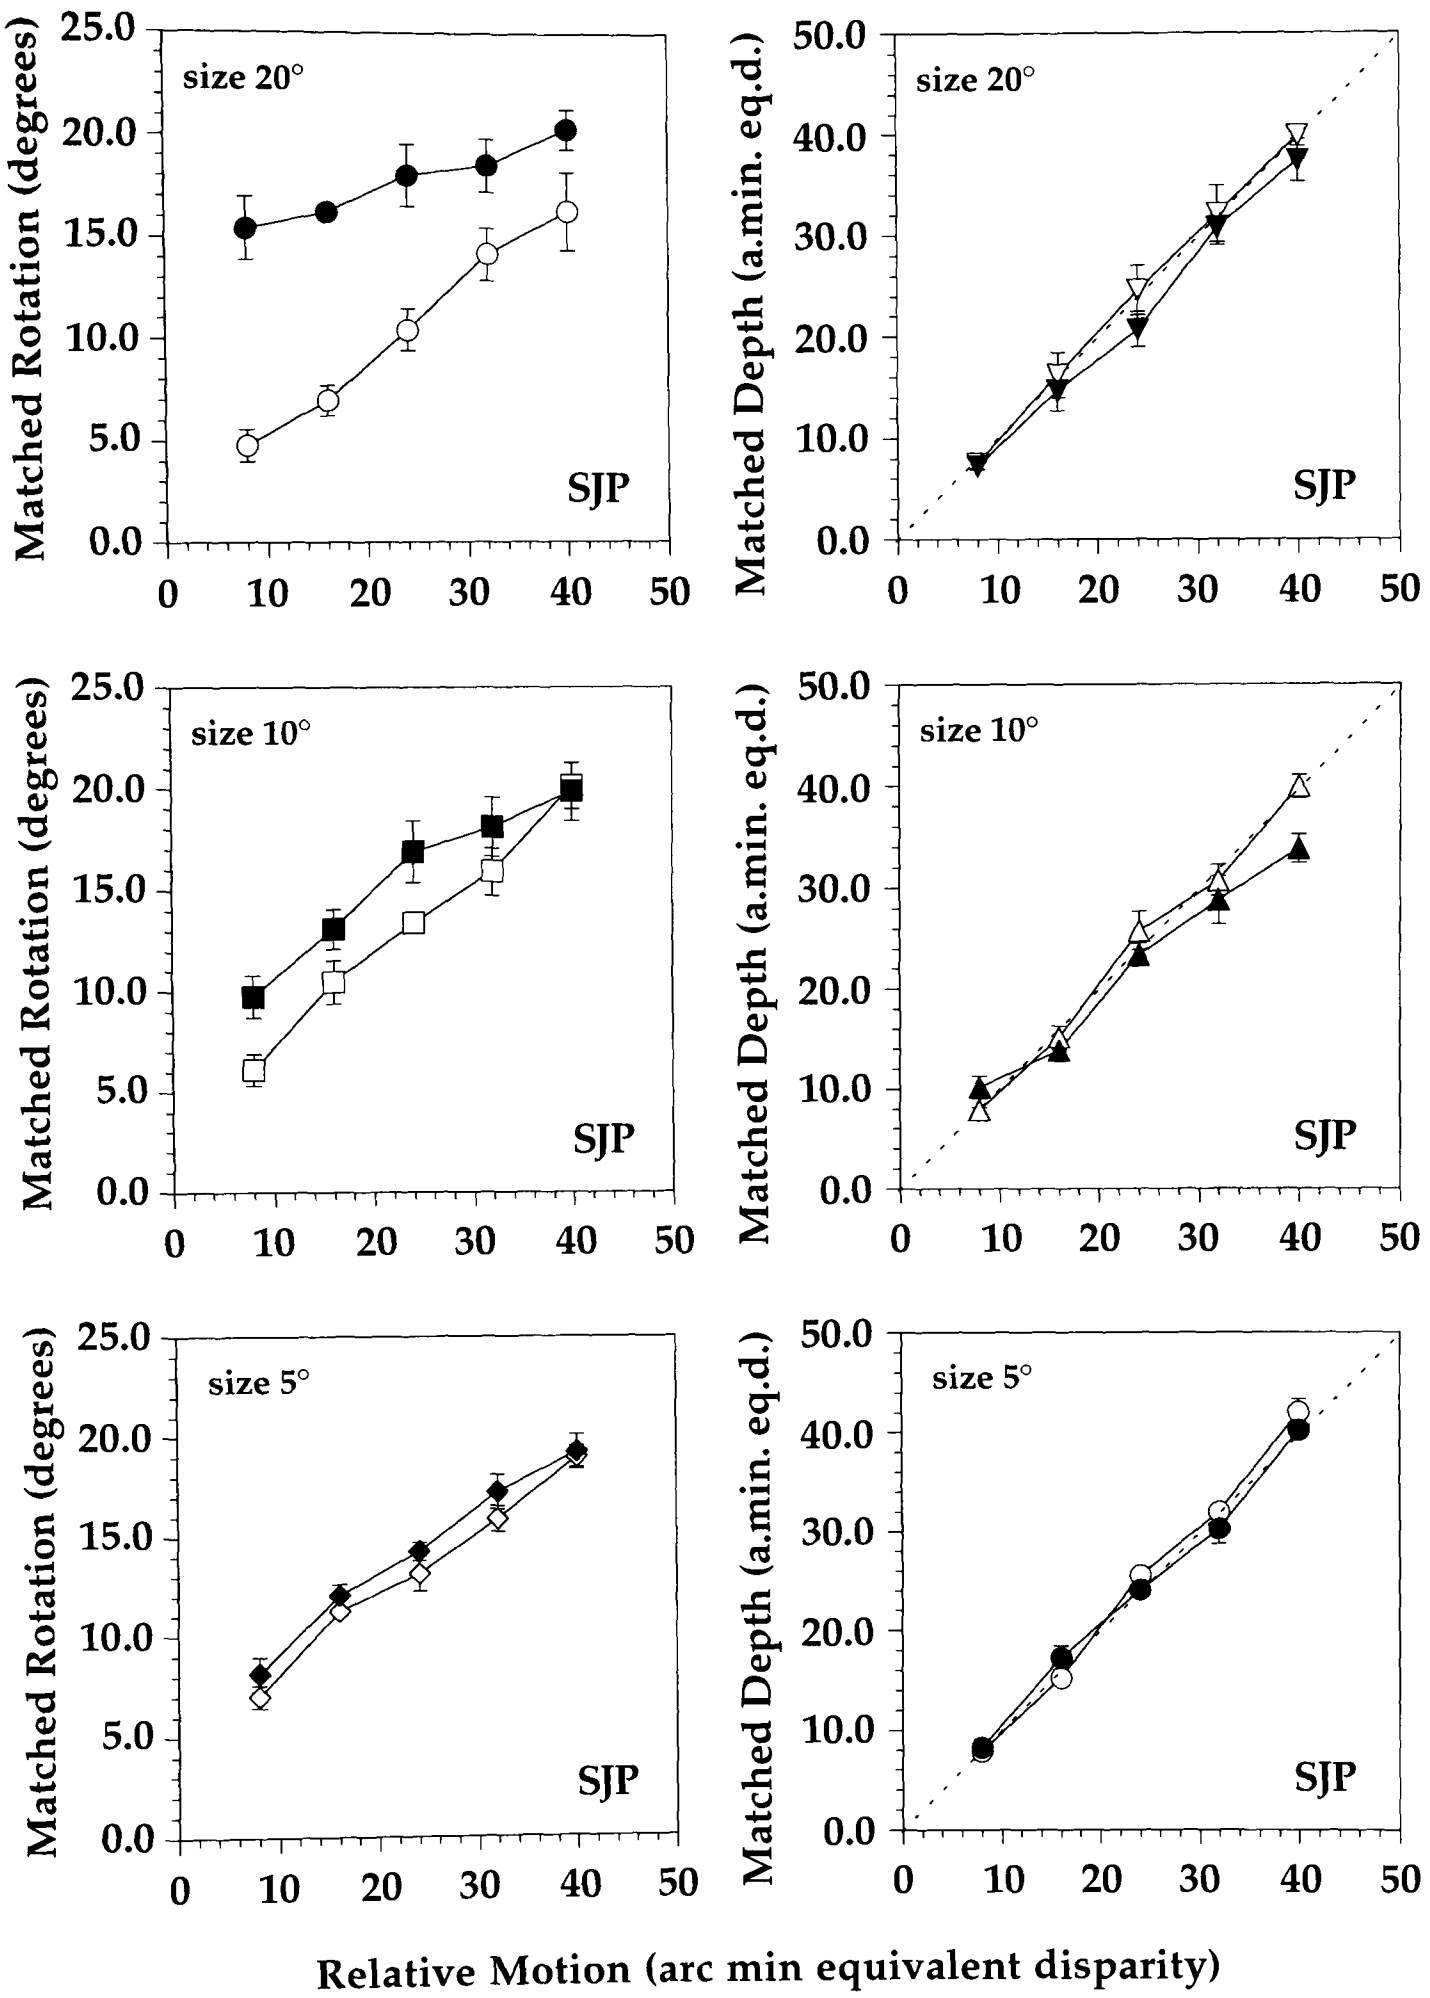


Fig. 6.9: Matched Rotation (left-hand panel) and Matched Depth (right-hand panel) for 0.4 cpd surfaces of size 5°, 10° and 20° for Subject SJP. The open symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 15°.

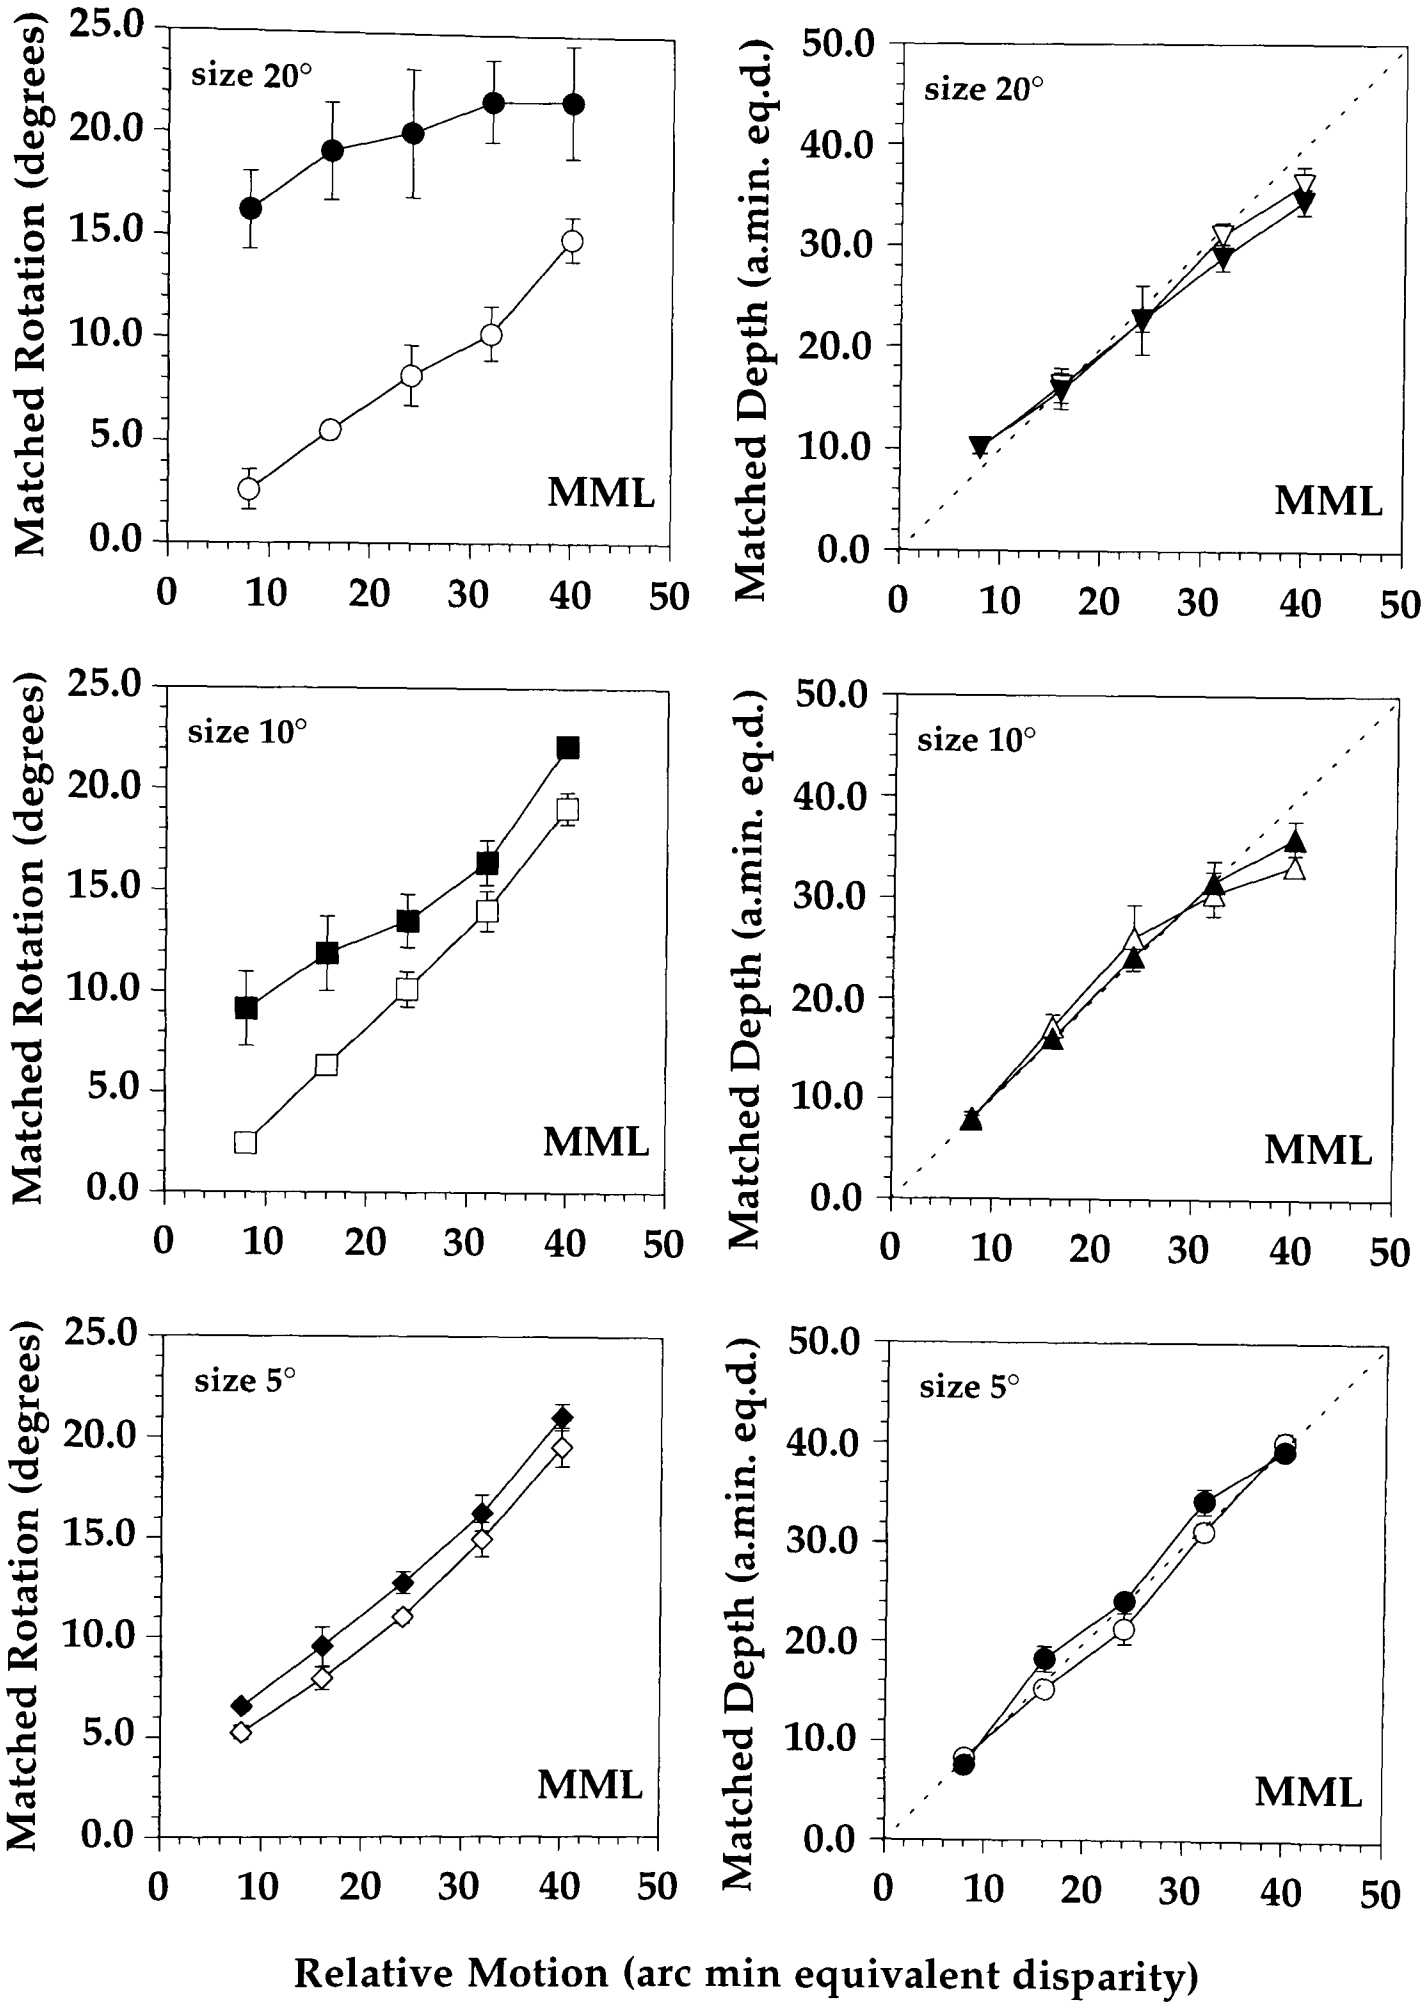


Fig. 6.10: Matched Rotation (left-hand panel) and Matched Depth (right-hand panel) for 0.4 cpd surfaces of size 5°, 10° and 20° for Subject MML. The open symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show Matches in the condition in which vertical perspective simulated a rotation angle of 15°.

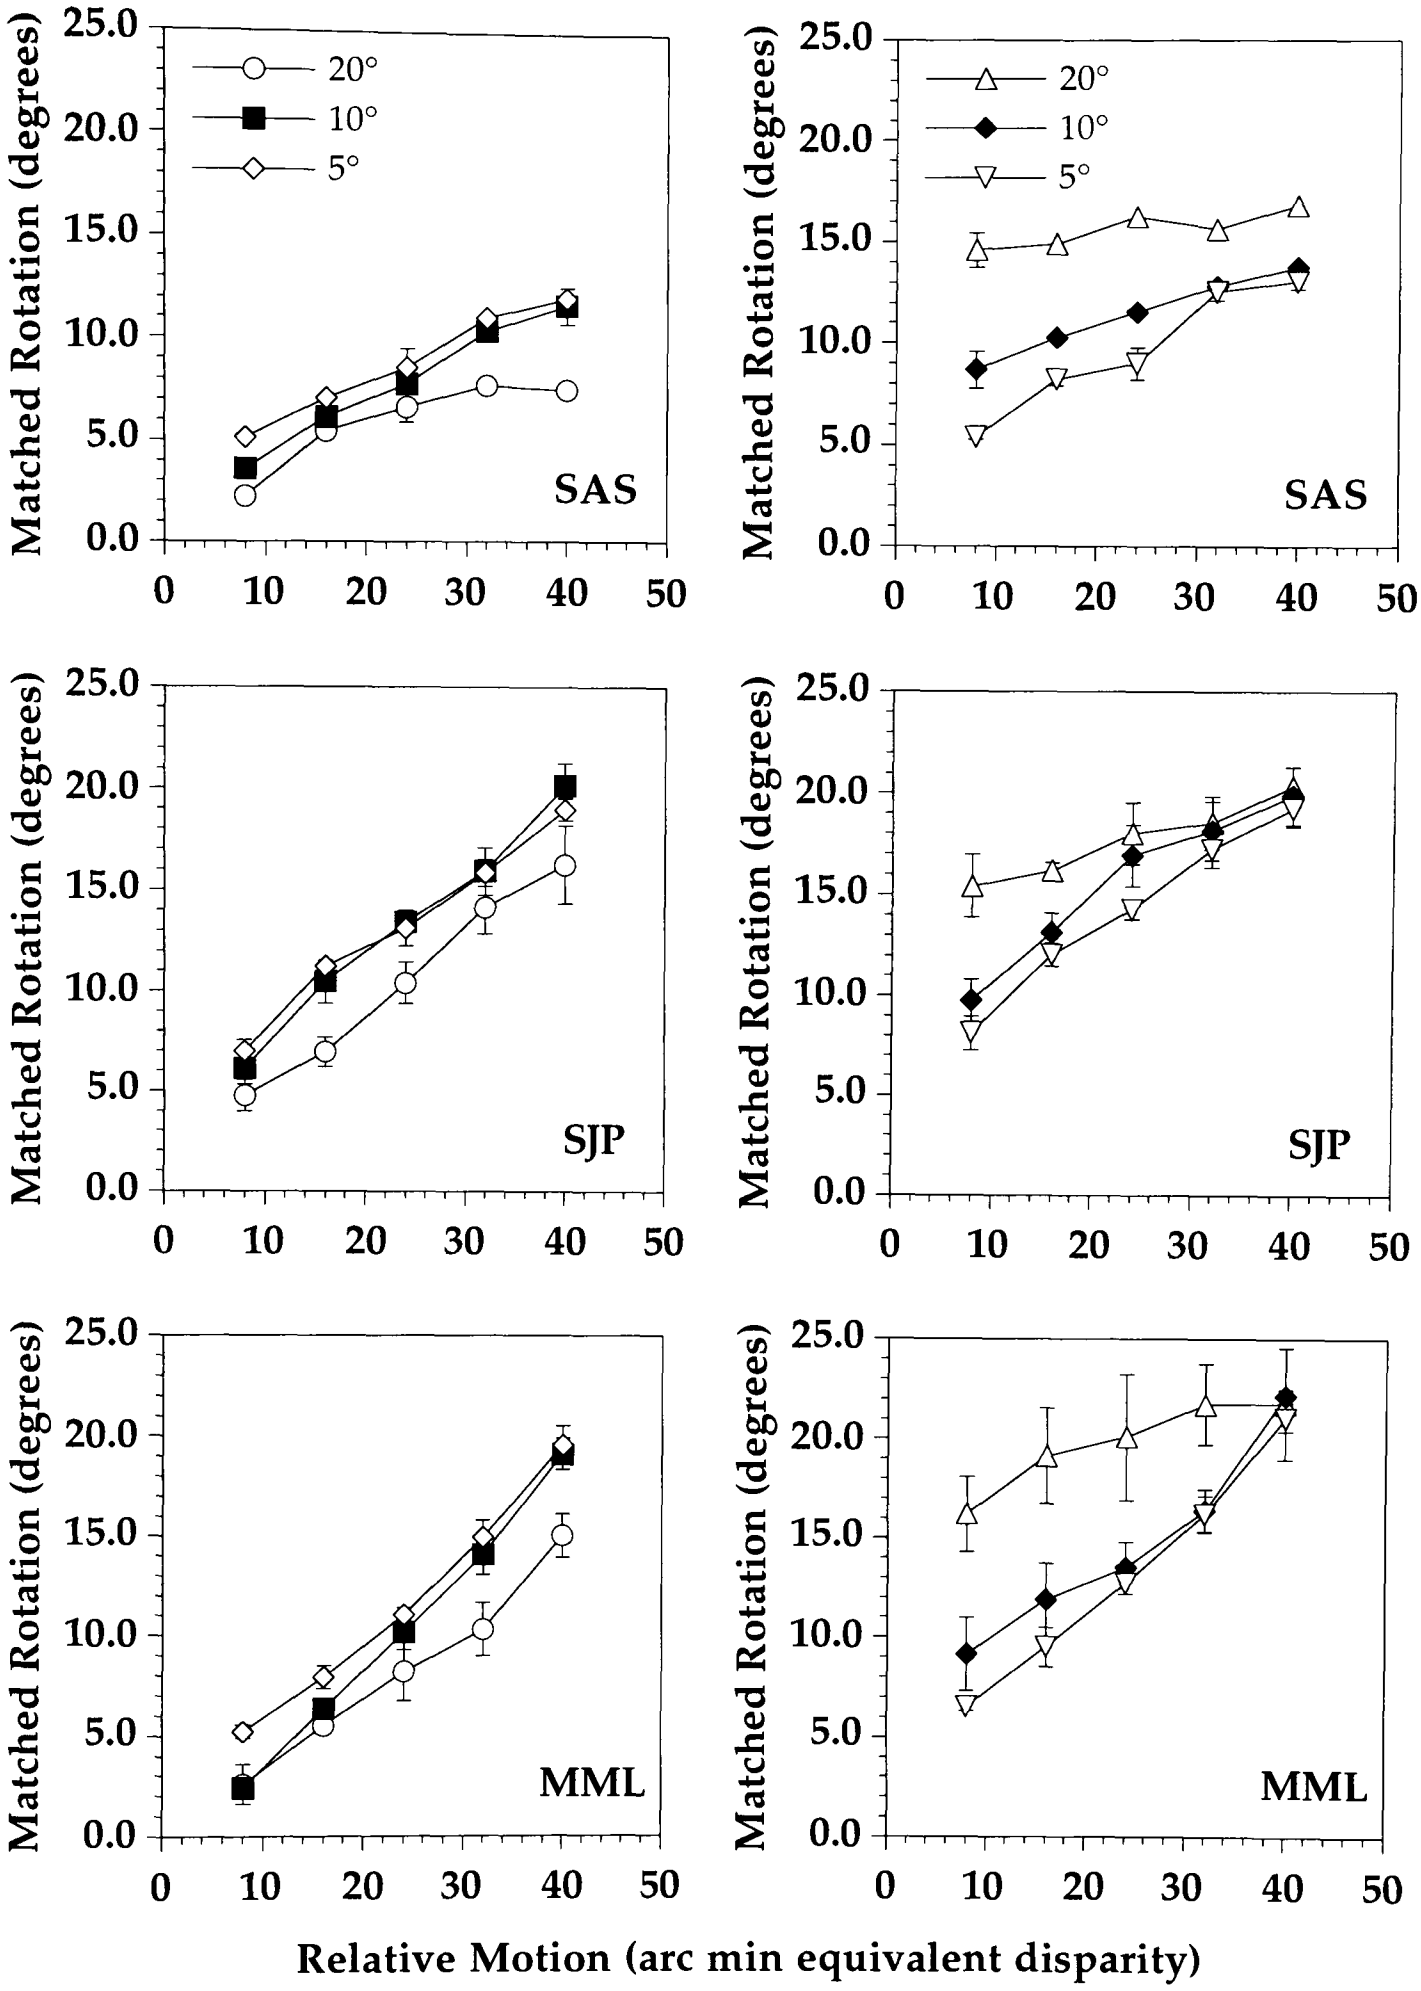


Fig 6.11: Matched Rotation for stimuli of size 5°, 10° and 20° for Subjects SAS, SJP and MML. The left-hand column shows matches for surfaces in which the magnitude of rotation signalled by the vertical perspective cue was 0°. The right-hand column shows matches for surfaces in which the magnitude of rotation signalled by the vertical perspective cue was 15°.

When the vertical perspective cue indicated a rotation angle of 15°, the

magnitude of perceived depth in the size 20° stimulus was slightly less than when the cue indicated an angle of 0°. This effect was either extremely small or absent for the size 10° and 5° stimuli.

Fig. 6.11 illustrates that when the vertical perspective cue signalled a rotation angle of 0° the magnitude of perceived rotation increased with decreasing stimulus size. When the perspective cue signalled a rotation angle of 15°, the magnitude of perceived rotation decreased with decreasing stimulus size.

6.3.4 Discussion

Dynamic vertical perspective was again shown to be a strong cue to the perceived rotation of observer-produced parallax surfaces. Its strength is dependent on the size of the stimulus. Thus when the signalled rotation was zero, the large stimulus appeared to rotate the least and the small stimulus was perceived to rotate the most. When the signalled rotation was 15°, this trend was reversed: the extent of perceived rotation was greatest for the largest stimulus and least for the smallest stimulus. This indicates that the effect of vertical perspective upon perceived rotation increases with increasing stimulus size. These results, although unsurprising, are novel since previous work quantifying the influence of vertical perspective on the magnitude of perceived rotation (e.g. Gibson and Gibson, 1957; Petersik, 1991; Williams, 1993) involved only stationary observers.

The effects on perceived rotation of vertical perspective and relative motion amplitude are not additive for the size 20° stimulus: the functions relating Matched Rotation to Relative Motion are shallower when the vertical perspective cue signals a non-zero rotation angle than when it signals a rotation angle of zero (i.e. is absent). This suggests that for this particular stimulus, under the conditions of the present experiment, the relative weightings of the perspective cue and the relative motion amplitude are qualitatively different according to whether the former signals a zero or a non-zero rotation angle. For smaller-sized stimuli, the relative weightings given to these two cues are relatively independent of the rotation angle signalled by the vertical perspective information.

The finding that slightly less depth was perceived in the majority of stimuli which were perceived to rotate more is qualitatively in line with the predictions of the trade-off model. However, as found in the previous experiment the large variations in perceived rotation which occurred with manipulation of the vertical perspective cue were not accompanied by correspondingly large variations in perceived depth.

6.4 General Discussion

The major conclusion which may be drawn from the experiments of this and the previous chapter is that when the magnitude of perceived rotation is manipulated through variation of a given stimulus attribute, the extent to which this is accompanied by a variation in the magnitude of perceived depth is dependent on the nature of the manipulated variable. Specifically, variations in corrugation frequency and in stimulus size led to variations in perceived rotation which were accompanied by substantial variations in perceived depth. However, variations in the angle of rotation simulated by the vertical perspective cue, while leading to variations in the magnitude of perceived rotation, generally did *not* lead to large variations in perceived depth. This is best illustrated by comparing the differences in perceived depth between pairs of stimuli containing approximately equal extents of relative motion and perceived rotation.

Experiment 11 showed that the perceived rotation of a size 20° stimulus with a relative motion amplitude of 36 arc min equivalent disparity was approximately 7° for a 0.05 cpd surface and 14° for a 0.4 cpd surface, for Subject SAS. The corresponding values for Subject SJP were 8° and 18°. These are shown in Figs. 6.12A and 6.13A below. For the former subject, the magnitudes of perceived depth were 34 and 64 mm for 0.4 and 0.05 cpd surfaces respectively (see Fig. 6.12B). The corresponding values for Subject SJP were 42 and 64 mm (Fig. 6.13B). In Experiment 16, the perceived rotation of a size 20° stimulus with a relative motion amplitude of 40 arc min equivalent disparity was approximately 7° for a 0.05 cpd surface in which the vertical perspective signalled a rotation of 0°, and 18° for the same surface when the perspective signalled a rotation of 15°, for Subject SAS (Fig. 6.12C). The corresponding values for Subject SJP were 8° and 16° (Fig. 6.13C). However, Figs. 6.12D and 6.13D show that despite the

significant difference in perceived rotation of the two surfaces, there was no difference between their perceived depths.

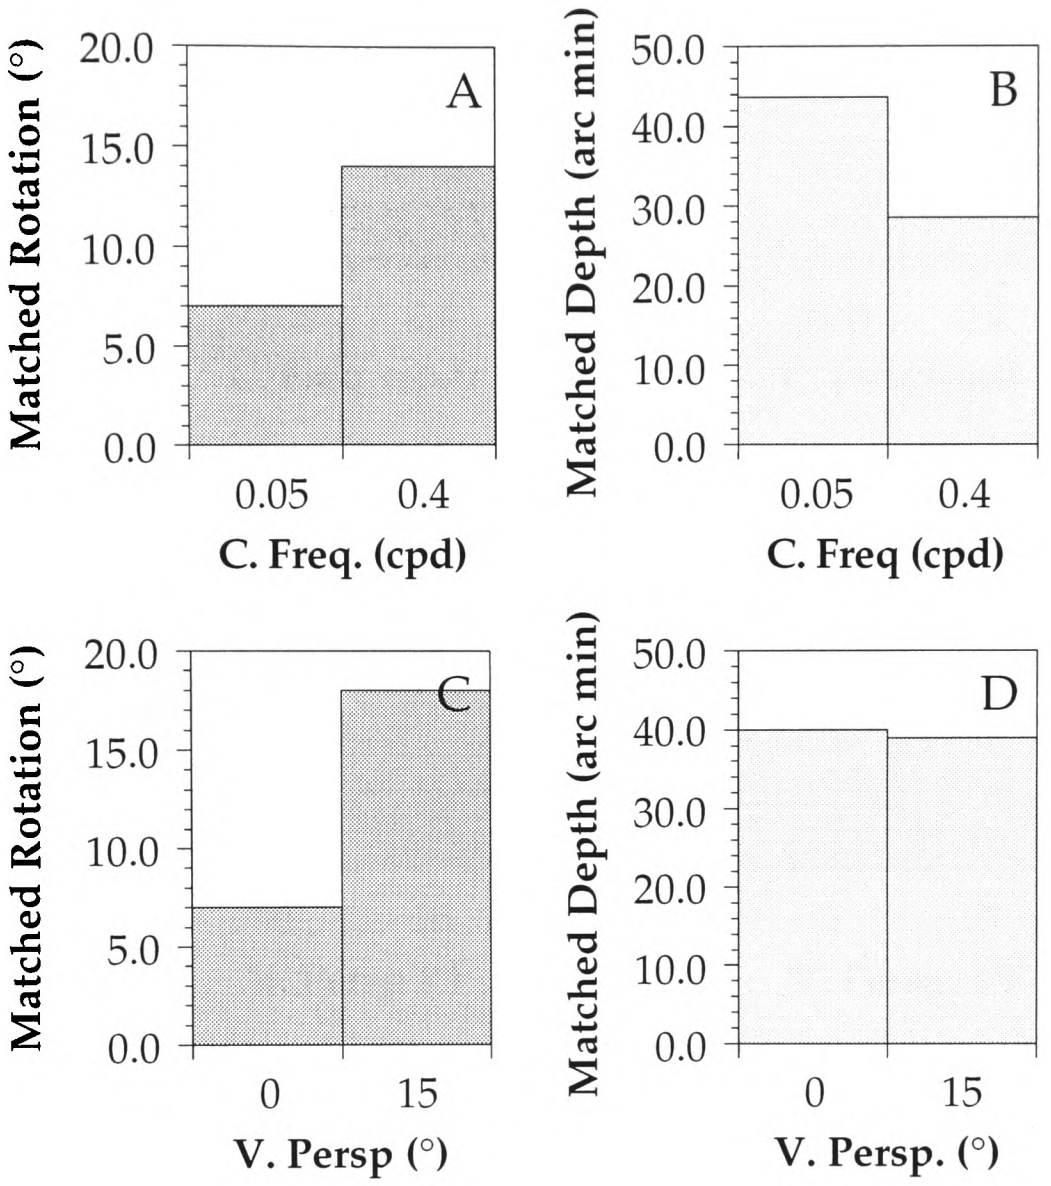


Fig 6.12: A & B - matched rotation and matched depth as a function of corrugation frequency, for Subject SAS (Experiment 11; 36 arc min equivalent disparity, size 20°). C & D - matched rotation and matched depth as a function of the rotation angle specified by the vertical perspective cue, for Subject SAS (Experiment 16; 40 arc min equivalent disparity, 0.05 cpd).

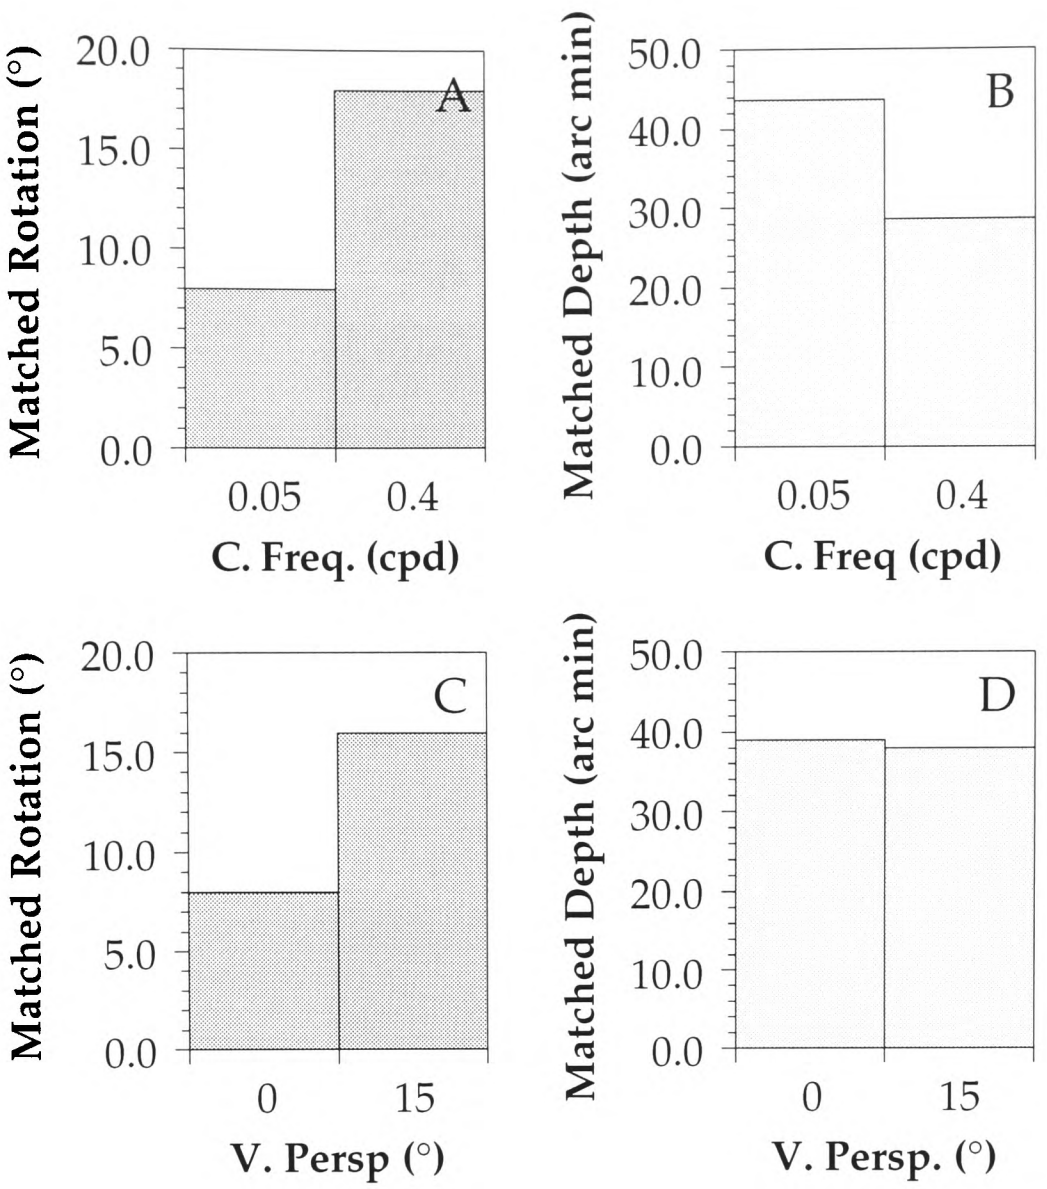


Fig. 6.13: A & B - matched rotation and matched depth as a function of corrugation frequency, for Subject SJP (Experiment 11; 36 arc min equivalent disparity, size 20°). C & D - matched rotation and matched depth as a function of the rotation angle specified by the vertical perspective cue, for Subject SJP (Experiment 16; 40 arc min equivalent disparity, 0.05 cpd).

Even in the conditions in which variations in vertical perspective *did* give rise to variations in the magnitudes of both perceived rotation and perceived depth, the size of the variations in the latter were generally either the same or smaller than those observed when either corrugation frequency or stimulus size was the manipulated variable. Fig. 6.14 compares the perceived depth and rotation between a stimulus from Experiment 12 and a stimulus from Experiment 17. The stimuli contained 20 and 24 arc min equivalent disparity respectively. The perceived rotation of the former stimulus was manipulated via stimulus size variations. The perceived rotation of the latter stimulus was manipulated through vertical perspective manipulations. It can be seen that although the effects of the given manipulation (stimulus size or vertical perspective) upon the perceived rotation magnitudes were similar, the effects upon the perceived depth magnitudes were not. Greater variations in perceived

depth resulted from the stimulus size manipulations than from the vertical perspective manipulations.

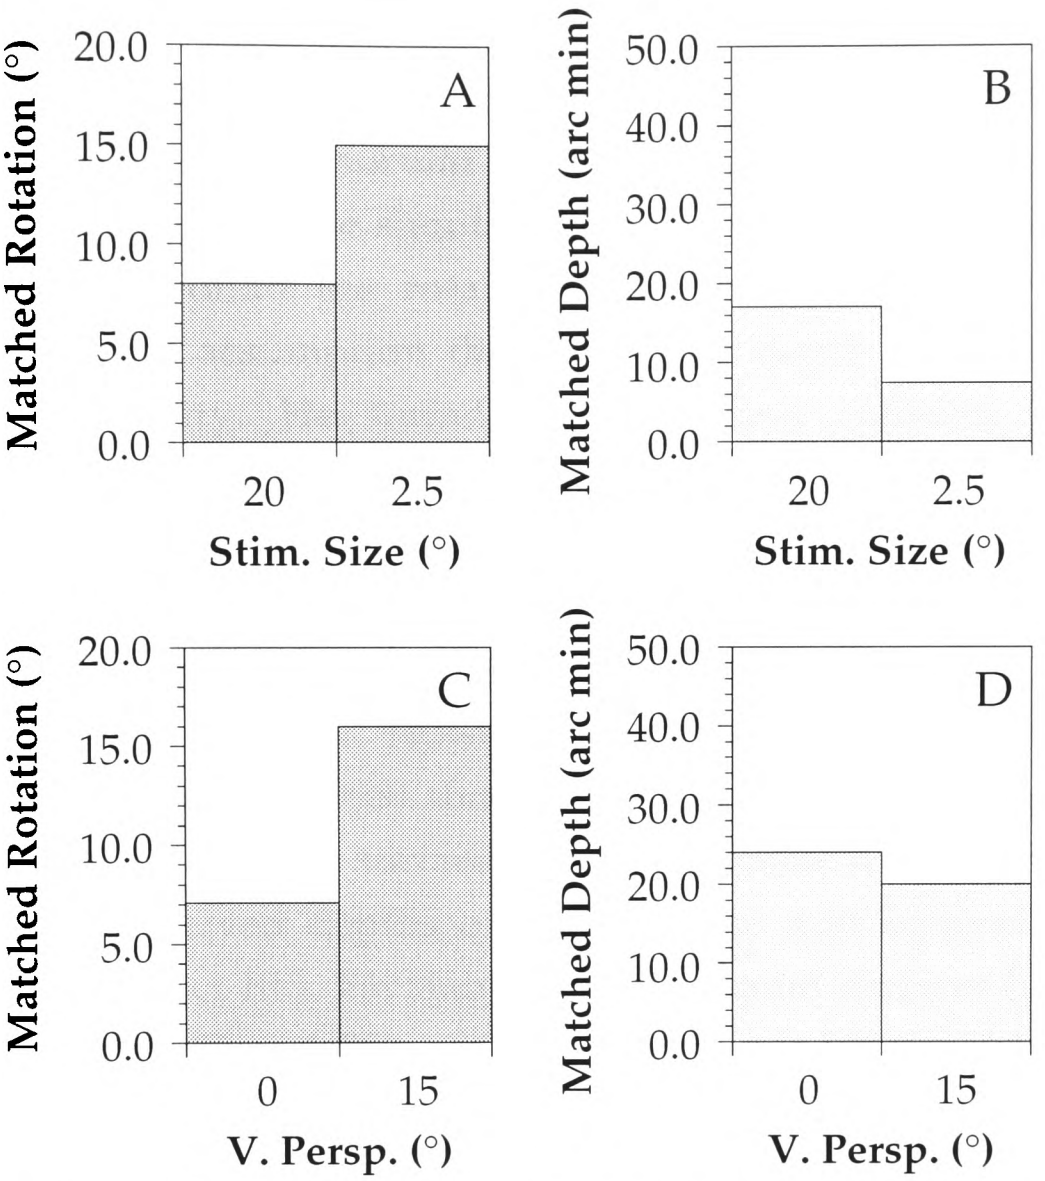


Fig 6.14: A & B - matched rotation and matched depth as a function of stimulus size, for Subject SAS (Experiment 12; 20 arc min equivalent disparity, 0.4 cpd). C & D - matched rotation and matched depth as a function of the rotation angle specified by the vertical perspective cue, for Subject SAS (Experiment 17; 24 arc min equivalent disparity, 0.4 cpd).

Although it is impossible to make direct comparisons between the depth matching data of Experiments 11 and 12 and those of the previous chapter because of the different techniques employed, the data strongly suggest that in cases where similar differences in perceived magnitudes of rotation are observed the differences between the extents of perceived depth are not necessarily similar. This in turn suggests that there is *not* a simple link between the extents of perceived rotation and depth such that a given amount of rotation necessarily coexists with a given amount of depth — but rather that the nature of the depth/rotation relationship is dependent upon the nature of the manipulated variable. Thus the trade-off relationship is observed when either the corrugation frequency or the

size of the stimulus are altered, but is observed to a significantly smaller degree when vertical perspective is the manipulated variable.

Williams (1993) investigated the role of dynamic vertical perspective in determining the perceived depth and rotation of surfaces defined both by motion parallax and by binocular disparity. He created a range of surfaces in which the rotation angle simulated by the vertical perspective cue took values centred around the rotation angle predicted by the geometric model, under the assumption that perceived depth was determined by binocular disparity. He found that matched rotation was largely determined by the vertical perspective cue, such that increases in the perspective level resulted in increases in the magnitude of perceived rotation. These increases were generally accompanied by decreases in the perceived depth magnitude, thus the data was interpreted as evidence for a depth/rotation trade-off. However, Figs. 7.6b, 7.7b and the graphs of Appendix G (Williams, 1993) show that the observed decreases in depth were small. For example, surfaces with perceived rotation angles of 10° and 45° had perceived depths of 12 and 8 arc min equivalent disparity respectively (Subject BL). For Subject JW, the same surfaces had perceived rotation angles of 0° and 58°, and had perceived depths of 10 and 7.5 arc min equivalent disparity respectively. These data were described by Williams (1993) as 'typical example cases' (p. 234). They illustrate that large increases in perceived rotation do not necessarily go hand in hand with large decreases in perceived depth, and thus bolster the notion that the variations in perceived depth which accompany variations in perceived rotation are dependent upon the nature of the manipulated variable(s).

The following two experiments investigated the putative trade-off between rotation and depth using two different shape judgement tasks.

6.5: Experiment 18: The Effect of Dynamic Vertical Perspective on Shape Judgements: Task One

6.5.1 Introduction

The present experiment involved a shape judgement task in which observers were asked to adjust the angle between the two planes of a triangular wave depth corrugation until the planes appeared to be at 90° to each other. If there exists a trade-off between the perceived magnitudes of rotation and depth, then it would be expected that the introduction of vertical perspective into a stimulus of fixed corrugation frequency would result in the observer requiring a *greater* amplitude of relative motion in order to make the structure appear right-angled than when no vertical perspective is present.

6.5.2 Stimuli

Surfaces contained horizontally-oriented, triangle wave depth corrugations. Corrugation frequency was either 0.15, 0.2 or 0.4 cpd. The magnitude of rotation simulated by the vertical perspective cue was either 0, +6.5°, +13°, -2° or -5°. Positively-signed angles indicate that the simulated rotation direction was opposite to the direction of head motion. Negatively-signed angles indicate that the simulated rotation direction was in the same direction as head motion. Observers controlled the peak-to-trough amplitude of the corrugations using a hand-held potentiometer. A small amplitude gave rise to a shallow corrugation, the sides of which met at a relatively large angle. A large amplitude produced a steeper corrugation, the sides of which met at a relatively small angle. Stimuli were 20° in diameter.

6.5.3 Procedure

Trials were blocked in groups of six, corresponding to two repetitions of each corrugation frequency. Subjects completed four experimental sessions, each consisting of five blocks, one at each level of vertical perspective. The order of trials in each block and the order of blocks in each session were randomised.

Observers were instructed to adjust the amplitude of relative motion until the sides of the central corrugation appeared to be at right angles to each other and thus resemble a (horizontally-oriented) chairleg. When subjects were satisfied with their 'right angle', vertical perspective was added to the stimulus, and a rotation match was made. Subjects then re-adjusted the relative motion amplitude until the stimulus appeared right-angled once more.

Between each adjustment of the relative motion amplitude the potentiometer was randomly set.

6.5.4 Results

Figs. 6.15 and 6.17 show the Rotation Match as a function of corrugation frequency for surfaces which had undergone the initial relative motion amplitude adjustment and had subsequently had vertical perspective added to them, for Subjects SAS and SJP respectively. The magnitude of perceived rotation increased with the rotation angle signalled by the vertical perspective cue.

Figs. 6.16 and 6.18 show the magnitude of relative motion required to make the stimulus appear right-angled, as a function of stimulus corrugation frequency. The different functions represent data obtained in each of the vertical perspective conditions. For a given corrugation frequency, an increase in the amount of rotation simulated by the vertical perspective cue led to an increase in the amount of relative motion required for a perceived right-angle — i.e. for a given relative motion amplitude, surfaces which appeared to rotate more appeared to contain less depth. It can also be seen that as corrugation frequency increased, the amount of relative motion required for the stimulus to appear right-angled decreased. The dashed lines show the relative motion amplitude theoretically required to produce a right-angle at each of the corrugation frequencies, under the assumption that the relative motion is interpreted as depth.

These graphs depict only the data for the conditions in which the angle of rotation simulated by the vertical perspective cue was 0, +6.5 or +13°. This is because the pattern of results obtained in these conditions is clear

for both subjects, whilst the data obtained in the conditions in which the simulated angle of rotation was -2 or -5° were less clear and their inclusion in the above graphs would result in a loss of clarity. The data of these latter conditions are presented separately in Figs. 6.19 - 6.22. Data of the 0° condition are also presented for comparison. It can be seen that both subjects perceived most rotation in the condition in which the simulated rotation was 0°, and least rotation in the condition in which the simulated rotation was -5°. The relative motion functions were not well separated, although both subjects required the greatest amount of relative motion in the condition in which the vertical perspective cue specified a rotation angle of 0°. These results are surprising at first sight, however they may be partially explained by the fact that neither subject perceived rotation in the direction simulated by the vertical perspective cue, a point which will be addressed in the following section.

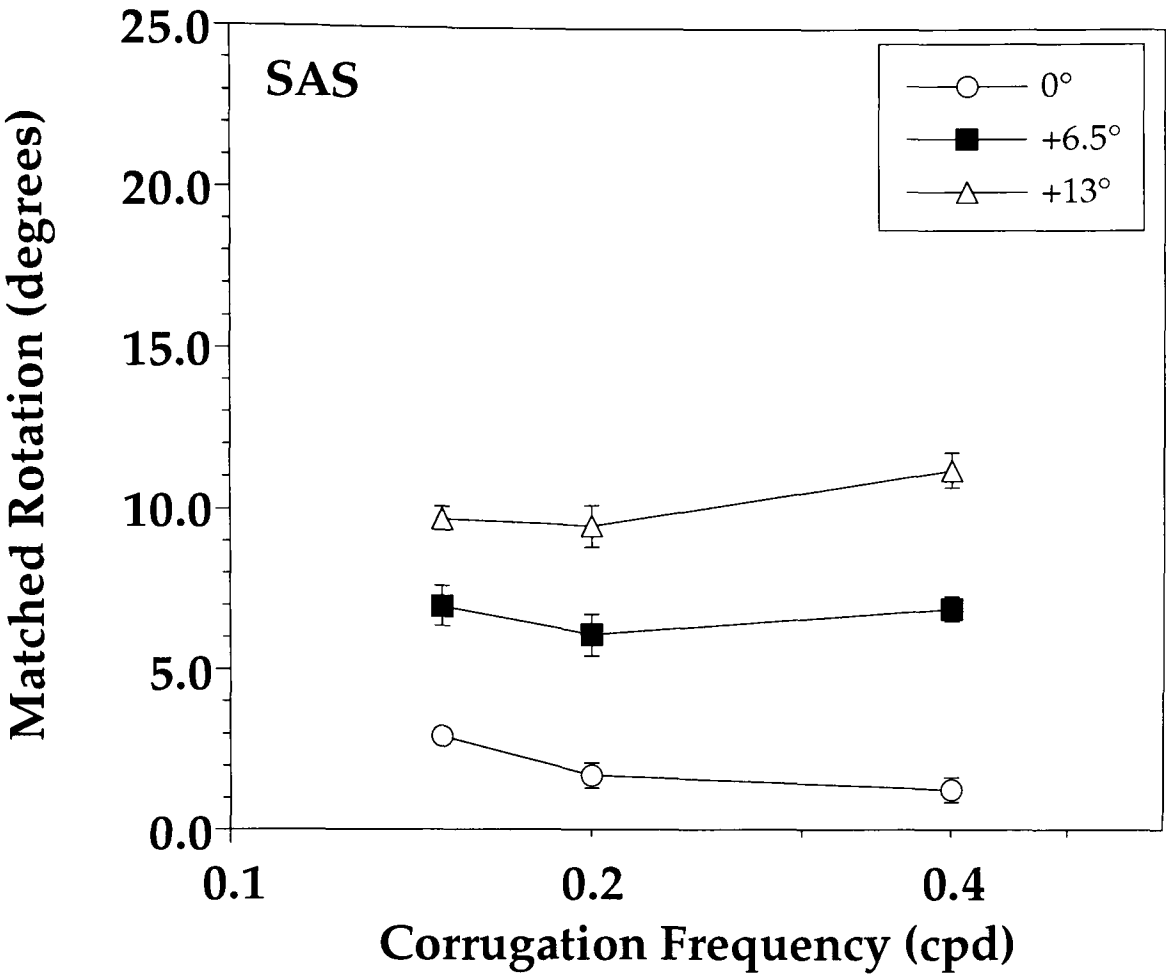


Fig. 6.15: Matched Rotation as a function of corrugation frequency for 'right-angled' stimuli for Subject SAS. The vertical perspective cue signalled a rotation angle of either 0°, +6.5° or +13°.

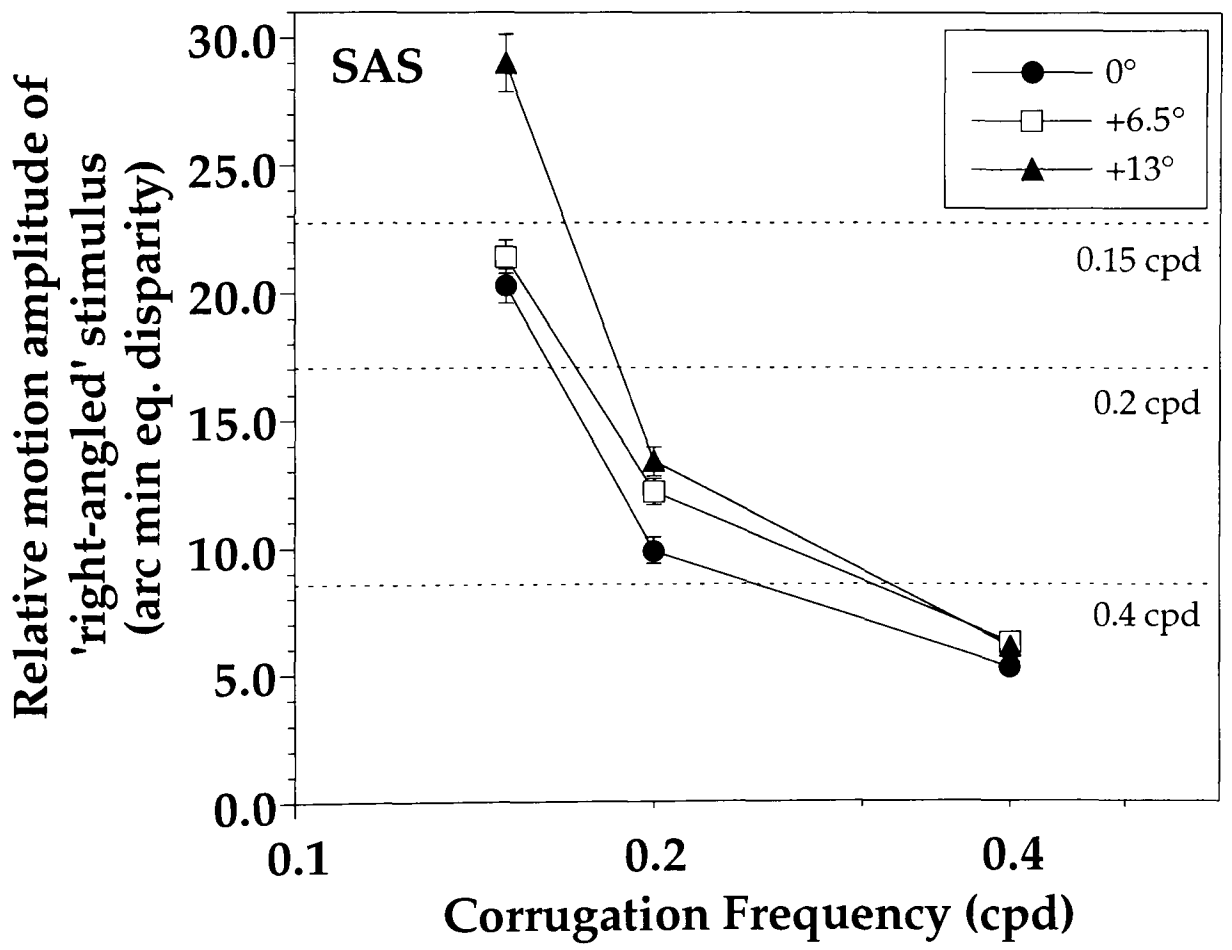


Fig. 6.16: Relative motion amplitude at which the stimulus appeared right-angled, as a function of corrugation frequency for Subject SAS. The rotation angle specified by the vertical perspective cue was either 0°, +6.5° or +13°. The dashed lines show the amplitude at which each corrugation frequency is physically right-angled.

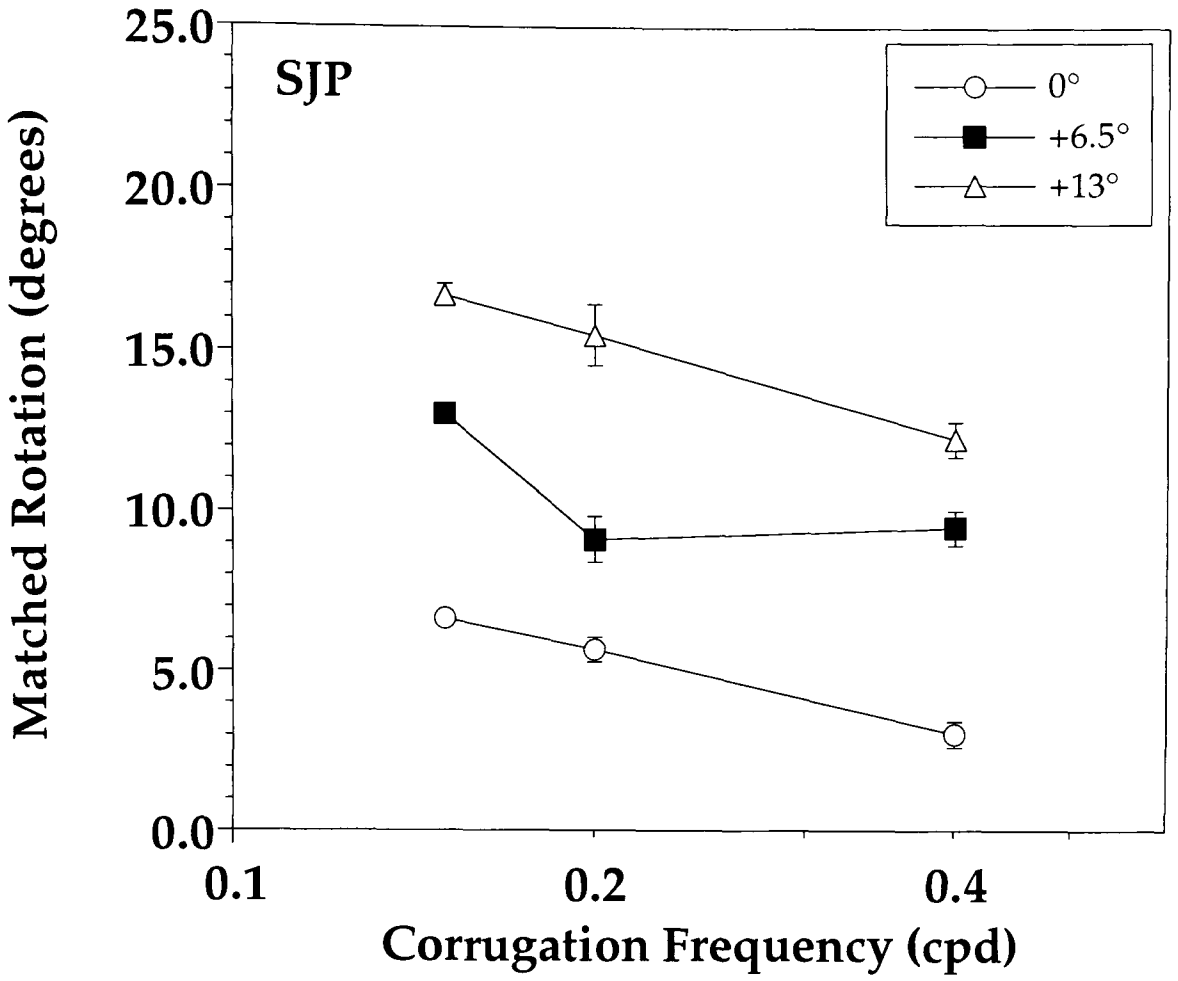


Fig. 6.17: Matched Rotation as a function of corrugation frequency for 'right-angled' stimuli for Subject SJP. The vertical perspective cue signalled a rotation angle of either 0°, +6.5° or +13°.

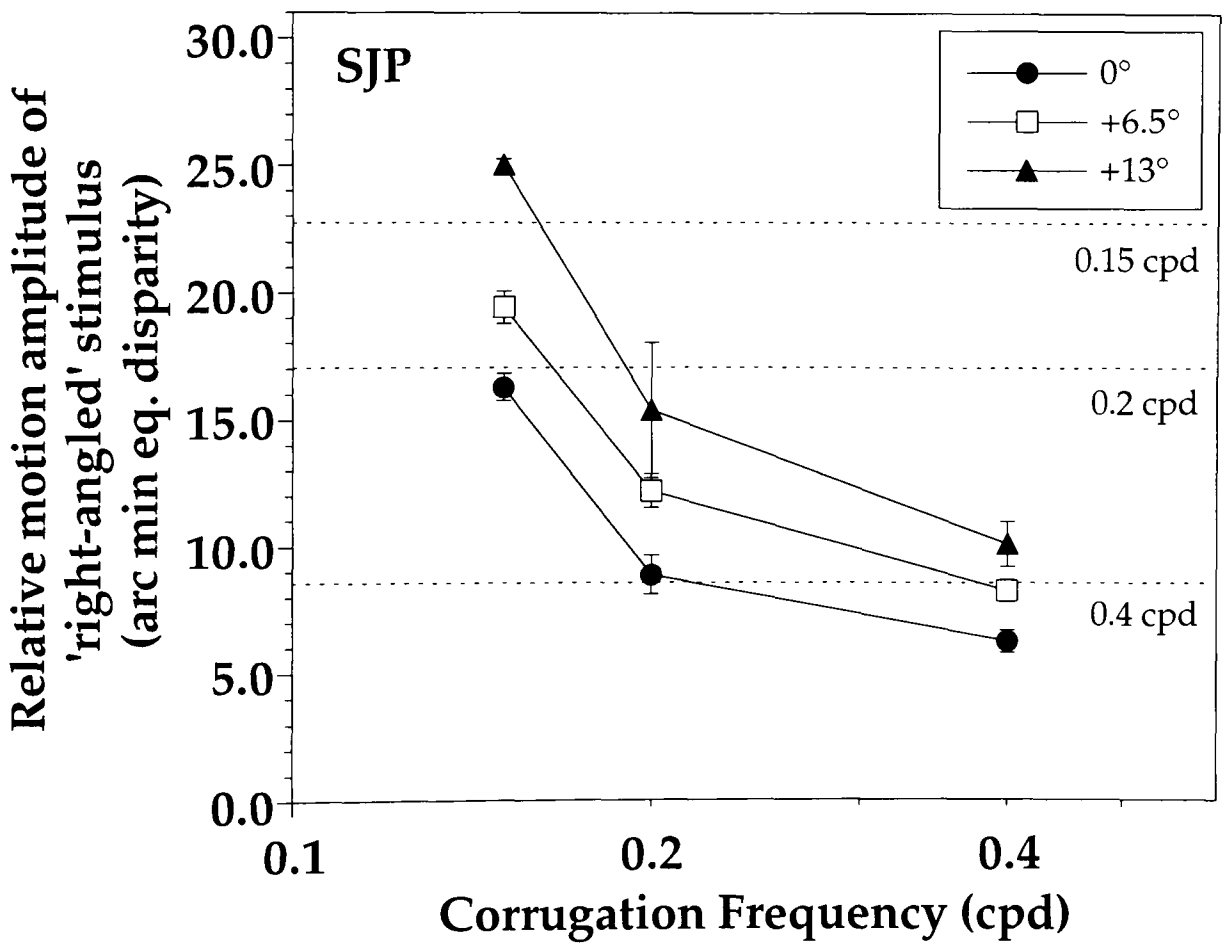


Fig. 6.18: Relative motion amplitude at which the stimulus appeared right-angled, as a function of corrugation frequency for Subject SAS. The rotation angle specified by the vertical perspective cue was either 0°, +6.5° or +13°. The dashed lines show the amplitude at which each corrugation frequency is physically right-angled.

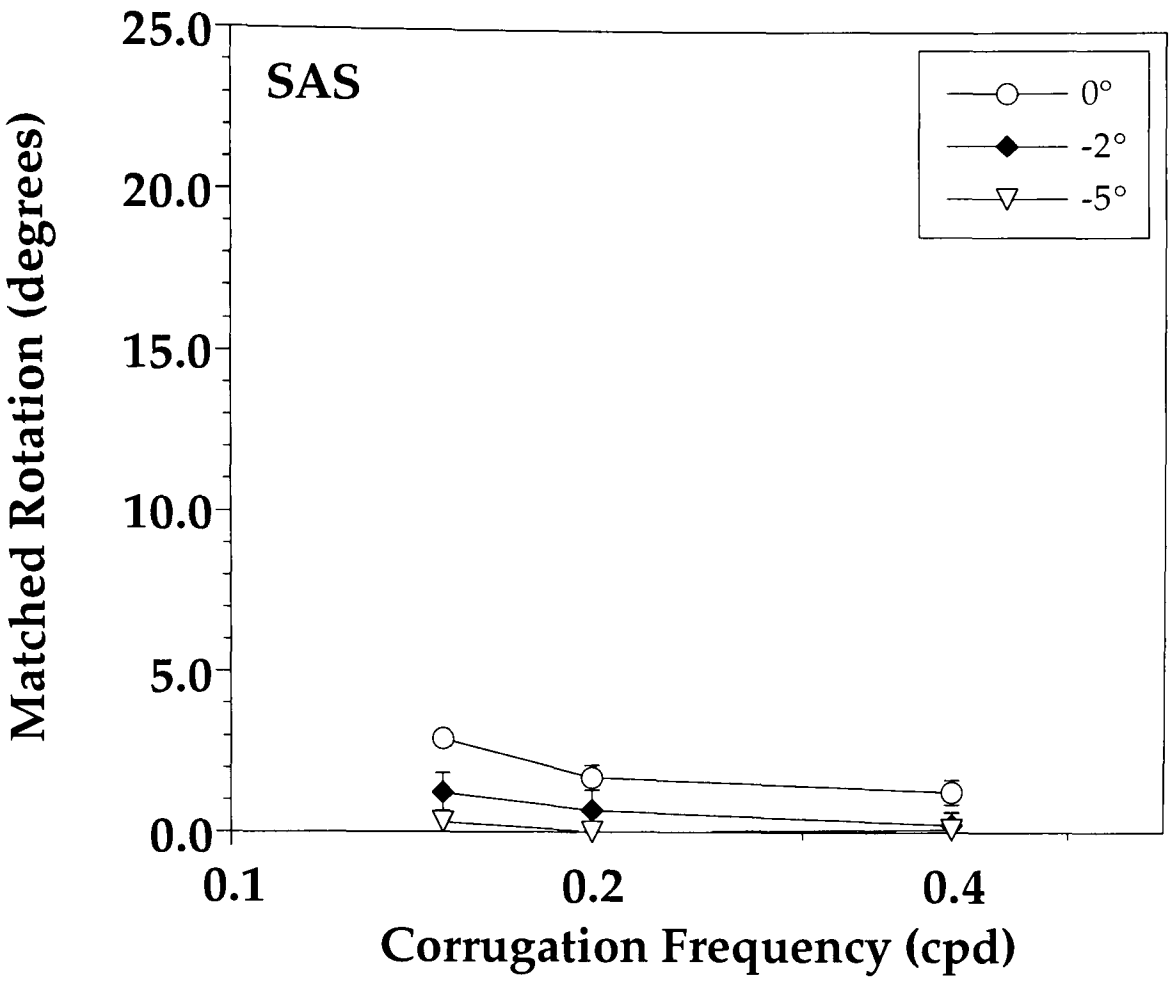


Fig. 6.19: Matched Rotation as a function of corrugation frequency for 'right-angled' stimuli for Subject SAS. The vertical perspective cue signalled a rotation angle of either 0°, -2° or -5°.

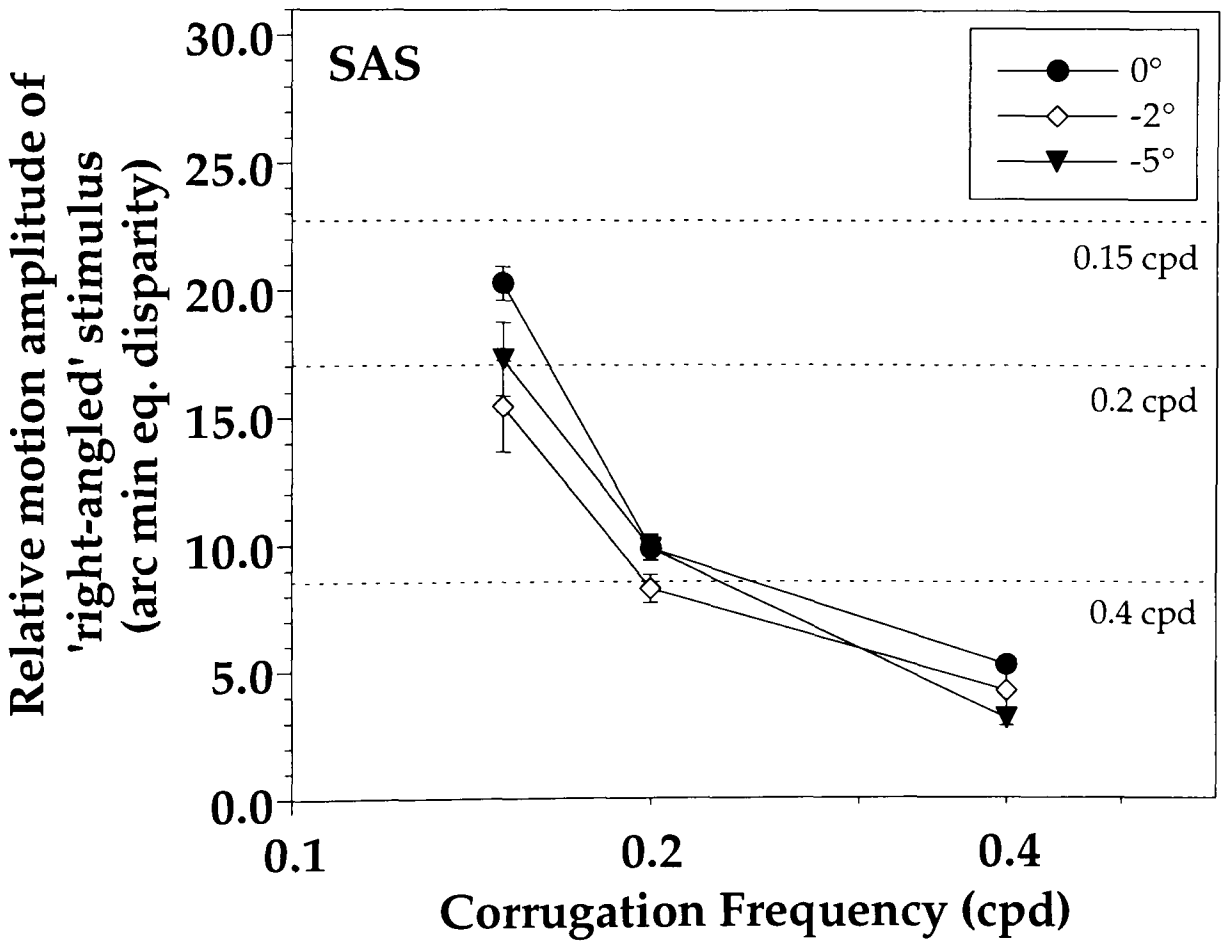


Fig. 6.20: Relative motion amplitude at which the stimulus appeared right-angled, as a function of corrugation frequency for Subject SAS. The rotation angle specified by the vertical perspective cue was either 0°, -2° or -5°. The dashed lines show the amplitude at which each corrugation frequency is physically right-angled.

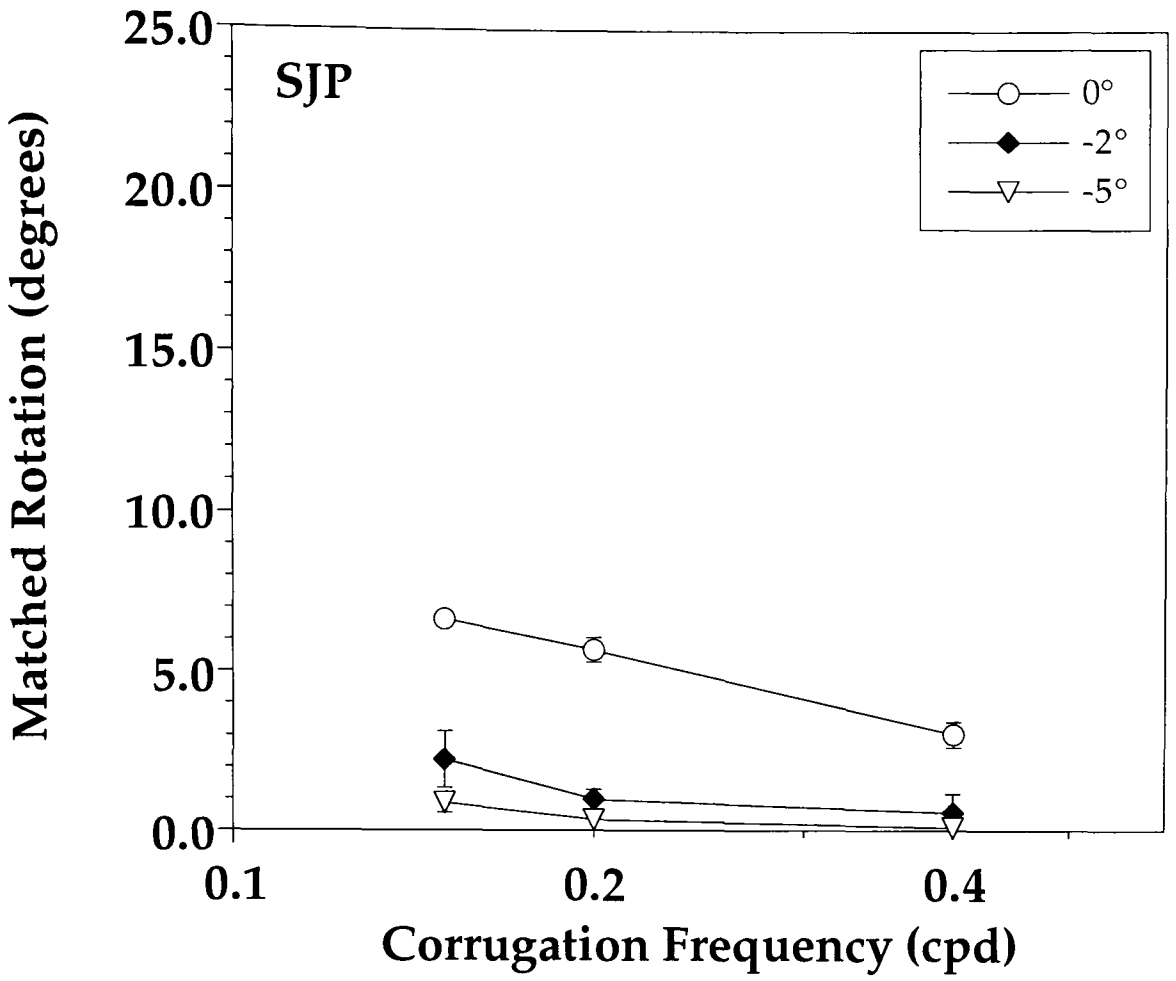


Fig. 6.21: Matched Rotation as a function of corrugation frequency for 'right-angled' stimuli for Subject SJP. The vertical perspective cue signalled a rotation angle of either 0°, -2° or -5°.

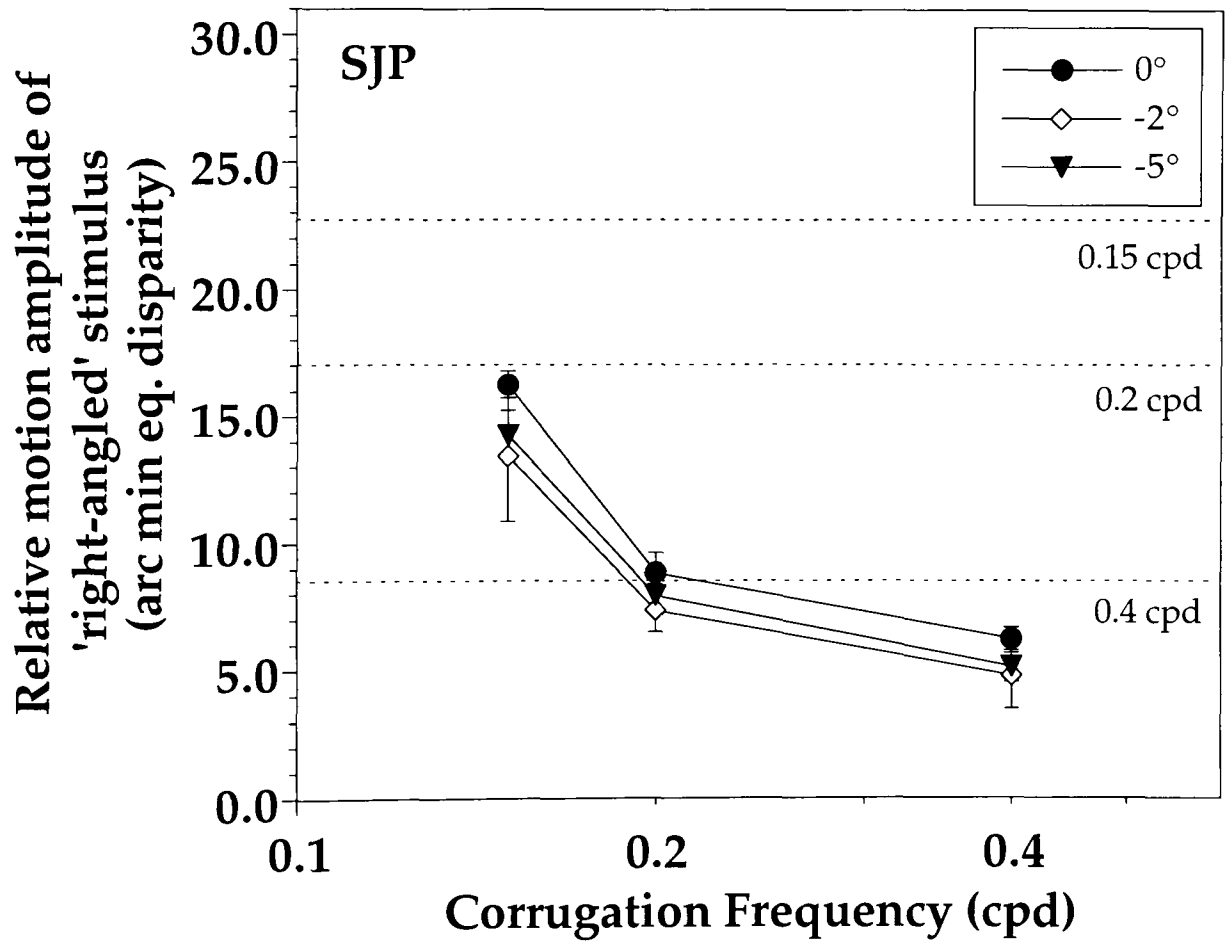


Fig. 6.22: Relative motion amplitude at which the stimulus appeared right-angled, as a function of corrugation frequency for Subject SJP. The rotation angle specified by the vertical perspective cue was either 0°, -2° or -5°. The dashed lines show the amplitude at which each corrugation frequency is physically right-angled.

6.5.5 Discussion

Considering first the results of the conditions in which the direction of rotation simulated by the vertical perspective cue was in the opposite direction to the head movement, it is evident that the amount of relative motion required to perceive a right-angled stimulus increased with increasing amounts of vertical perspective: surfaces which were perceived to rotate to a greater extent were perceived to contain less depth. This finding is in accord with the notion that there is a trade-off between perceived depth and motion. As the perceived rotation of the stimulus increases, the portion of relative motion interpreted as arising from the depth of the corrugations decreases. Consequently the angle appears to be greater and therefore more relative motion is required in order to make it appear to be 90°. The effect is small, but becomes clearer as the corrugation frequency decreases.

The absolute relative motion amplitudes required to perceive a right-angle are considerably smaller than predicted. Similar under-estimations of the amount of depth required for the perception of a given shape have been found in other studies (e.g. Johnston, 1991; Glennerster, Rogers and Bradshaw, 1996; Tittle, Todd, Perotti and Norman, 1995). Johnston (1991) presented observers with random dot stereograms depicting a horizontally-oriented elliptical profile. The task was to adjust the amplitude so that the profile appeared semi-cylindrical (i.e. so that the depth was equal to half the height, which took one of five values between 2.5 and 7.5 cm). At a viewing distance of 53.5 cm the apparently circular cylinder (ACC) corresponded to a physically elliptical cylinder flattened along its depth axis, indicating that the perceived amount of depth was over-estimated. At the two greater viewing distances used in this study (107 and 214 cm) the amount of depth was under-estimated: the ACC corresponded to a physically elliptical cylinder stretched along its depth axis. Glennerster et al. (1996) asked subjects to perform a task essentially the same as that of Johnston (1991). Using elliptical hemi-cylinders with half-heights of either 3 or 6 cm, subjects again varied the amplitude until the profile appeared semi-circular. For viewing distances of 38, 57 and 76 cm the depth of the surface which appeared cylindrical was less than the half-height of the profile. For the viewing distances of 114 and 228 cm this trend was reversed and a greater depth than the half-height was required.

In an additional experiment observers adjusted the amplitude of a triangle wave until the dihedral angle appeared to be 90°. At all viewing distances except 228 cm depth was over-estimated and observers required a smaller peak-to-trough height than predicted.

Tittle et al. (1995) obtained similar results to those above when they carried out the elliptical hemi-cylinder task using stereoscopic stimuli. They also carried out the task using perspective projections of horizontally oriented cylinders rotating ±5° about an average slant of 0, 15 or 30° out of the frontal plane. Depth over-estimation occurred when the average slant was 0°. In a second task observers were required to adjust the dihedral angle of a triangle wave so that it appeared to be 90°. Once more the average slant out of the frontal plane was 0, 15 or 30°. For stereoscopic stimuli there was an over-estimation of depth for the nearest viewing distance, but there was no effect of slant. For the parallax stimuli there was a large degree of depth over-estimation when the average slant was 0°. As in the cylinder adjustment task there was little effect of viewing distance on perceived depth.

These shape judgement tasks indicate that there is a degree of depth over-estimation at near viewing distances, both for parallax-defined and stereoscopic surfaces. The depth over-estimation of stereoscopic surfaces has often been attributed to an over-estimation of the viewing distance. The depth over-estimation of motion parallax stimuli could arise from a number of factors, which were discussed in Section 5.4.

Perceived depth order and direction of rotation

Stimuli in which the vertical perspective cue simulated rotation in the same direction as the head movement were *not* perceived to rotate in the simulated direction, but rather were perceived to counter-rotate and/or to appear non-rigid. Both subjects perceived least rotation in the stimulus in which the magnitude of simulated rotation was -5°, and the greatest amount of rotation in the stimulus in which the magnitude of simulated rotation was 0°. Both subjects perceived non-rigidity in the former stimulus, although quantitative measurements were not made. It is clear that vertical perspective, which has been shown to be a strong cue to the direction and magnitude of rotation (e.g. Gibson and Gibson, 1957; Power,

1967; Braunstein and Payne, 1968; Börjesson, 1971; Hershberger et al., 1976; Braunstein, 1977; Rogers and Rogers, 1992) does not dominate the rotation percept in the present circumstances.

Inspection of the depth/rotation combinations allowed by the geometric model (Fig. 5.6) illustrates that, for a given relative motion amplitude, the rigid surface can theoretically be perceived as rotating either in the same direction as the head or in the opposite direction to the head. The depth order when rotation is in the latter direction is necessarily positive. That is, dots moving in the opposite direction to the head are perceived as lying on peaks and dots moving in the same direction as the head are perceived as lying on troughs. The depth order when rotation is in the former direction is positive when the angle of rotation lies between -6.5 and 0°, and is negative — dots moving in the opposite direction to the head are perceived as lying on troughs and dots moving in the same direction as the head are perceived as lying on peaks — when the angle of rotation is less than -6.5°.

Despite the fact that the geometrically permissible solution for a given amount of relative motion may have either of two opposite depth orders, the perceived order in head-linked motion parallax surfaces has been reported without exception as positive (e.g. Rogers and Graham, 1979; Ono et al., 1986; Rivest et al., 1989; Ono and Steinbach, 1990). The propensity of the visual system to perceive a positive depth order may arise from a general bias to perceive elements moving in the opposite direction to the head as lying in front of those moving in the same direction as the head. This is a sensible heuristic given that the retinal images of stationary objects in the real world behave in this way. The bias may be further strengthened by the fact that the magnitude of rotation associated with a positive depth order can be small (and in either direction, as illustrated in Fig. 5.6), whereas the magnitude of (negative) rotation associated with a negative depth order must be more negative than -6.5°. A further contributing factor to this bias is likely to be the (real-world) vertical perspective information arising from the display upon observer movement. This information specifies a stationary, frontal surface. With respect to the eye, however, the surface rotates in the opposite direction to the head movement. A negative depth order is geometrically inconsistent with this rotation direction (see Fig. 5.6).

The angle of rotation simulated by the vertical perspective cue in the current experiment was 0, +6.5, +13, -2 or -5°. In a pilot study, angles of greater simulated negative magnitude were found to produce surfaces of positive perceived depth order which were consequently extremely non-rigid in appearance and which did not appear to rotate. This underlines the propensity of the visual system to perceive a positive depth order, even at the expense of rigidity. It does not, however, indicate whether this interpretation of the stimulus is due to a preference to perceive a positive depth order *per se* or whether it is due to a reluctance to perceive rotation in the same direction as head movement. This question was addressed by the stimuli in which the simulated magnitude of rotation was either -2 or -5°. In these cases, the perceived depth order was positive (consistent with the predictions of the model) and the perceived rotation direction was also positive. This result suggests that the visual system is biased against perceiving rotation in the same direction as head motion, even when this does *not* necessitate the perception of a negative depth order. It therefore suggests that observers have a bias for interpreting elements which move in the opposite direction to head motion as counter-rotating which is independent of the bias to perceive a positive depth order.

To summarise, the results of the current experiment show that the relative motion amplitude required to produce an apparently right-angled structure increases with increases in the magnitude of rotation simulated by the vertical perspective cue. The absolute amounts of relative motion required for the task are considerably less than predicted. This is in keeping with shape judgement tasks in the literature, and with the results of Experiments 11 - 15 which showed that the combined extents of perceived rotation (or shearing) and depth were greater than those predicted by the trade-off model.

6.6: Experiment 19: The Effect of Dynamic Vertical Perspective on Shape Judgements: Task Two

6.6.1 Introduction

The previous experiment demonstrated that a greater peak-to-trough

depth amplitude was required to make the surface appear right-angled when vertical perspective was present than when it was absent. The present experiment used a different stimulus manipulation to approach the depth/rotation trade-off question. In a surface in which the depth amplitude is fixed and observers adjust the corrugation frequency, it would be expected that the frequency at which the stimulus appears right-angled would be greater in the presence of vertical perspective than in a condition in which it was absent.

6.6.2 Stimuli

Stimuli were identical to those of the previous experiment, except that the Peak-to-Trough relative motion amplitude was now fixed and the corrugation frequency was adjustable. The relative motion amplitude was 10, 20, 30 or 40 arc min equivalent disparity.

6.6.3 Procedure

Trials were blocked in groups of eight, corresponding to two repetitions of each relative motion amplitude. Subjects completed four experimental sessions, each consisting of five blocks, one at each level of vertical perspective. The order of trials in each block and the order of blocks in each session were randomised.

The procedure was identical to the previous experiment, except that the hand-held potentiometer controlled the corrugation frequency of the stimuli.

6.6.4 Results

Figs. 6.23 - 6.26 show data for the conditions in which the rotation angle simulated by the vertical perspective cue was 0, +6.5 or +13°. Figs. 6.23 and 6.25 show Matched Rotation as a function of relative motion amplitude for surfaces the corrugation frequency of which had been adjusted to create a perceived right-angle, and to which the vertical perspective cue had subsequently been added, for subjects SAS and SJP respectively. The separate functions represent the different vertical perspective conditions. The data show the anticipated pattern: the perceived magnitude of

rotation increases with increasing amounts of vertical perspective. Figs. 6.24 and 6.26 show the corrugation frequency required to make the stimulus appear right-angled as a function of relative motion amplitude, for each of the three rotation angles simulated by the vertical perspective cue. It can be seen that for a given amplitude of relative motion, as the simulated rotation angle increases so does the corrugation frequency required to produce a 'right-angled' surface.

Figs. 6.27 - 6.30 show the data from the conditions in which the vertical perspective cue simulated a rotation angle of 0, -2 or -5°. There is no clear pattern of results, either for the perceived rotation data or for the corrugation frequency adjustment data.

6.6.5 Discussion

When the vertical perspective cue signalled a rotation in the opposite direction to the head movement, increases in the simulated rotation angle led to increases in both the magnitude of perceived stimulus rotation and the corrugation frequency at which the stimulus appeared right-angled. Once again these results are consistent with a trade-off between depth and rotation. It can also be seen that the corrugation frequency at which the stimulus appeared right-angled was lower than predicted, indicating that there was an over-estimation of the amount of depth in the stimulus. The results were again unclear for the surfaces in which the angle of rotation specified by the vertical perspective cue was negative. Once again these stimuli were perceived to rotate against, rather than with, the direction of head movement, and were perceived as non-rigid to varying degrees.

Overall, the results from the two experiments indicate that there is a small trade-off between perceived rotation and perceived depth, as found in Experiments 11 - 16. The results also indicate that the amount of depth in the surfaces is over-estimated, a finding which is consistent with previous work involving both stereoscopic (Glennerster et al., 1996; Johnston, 1991) and parallax-defined (Durgin et al., 1995) surfaces. Additionally, it was demonstrated that observers do not readily perceive rotation which is in the same direction as that of the head motion. This bias is independent of whether the 'correct' solution involves a positive or a negative depth order.

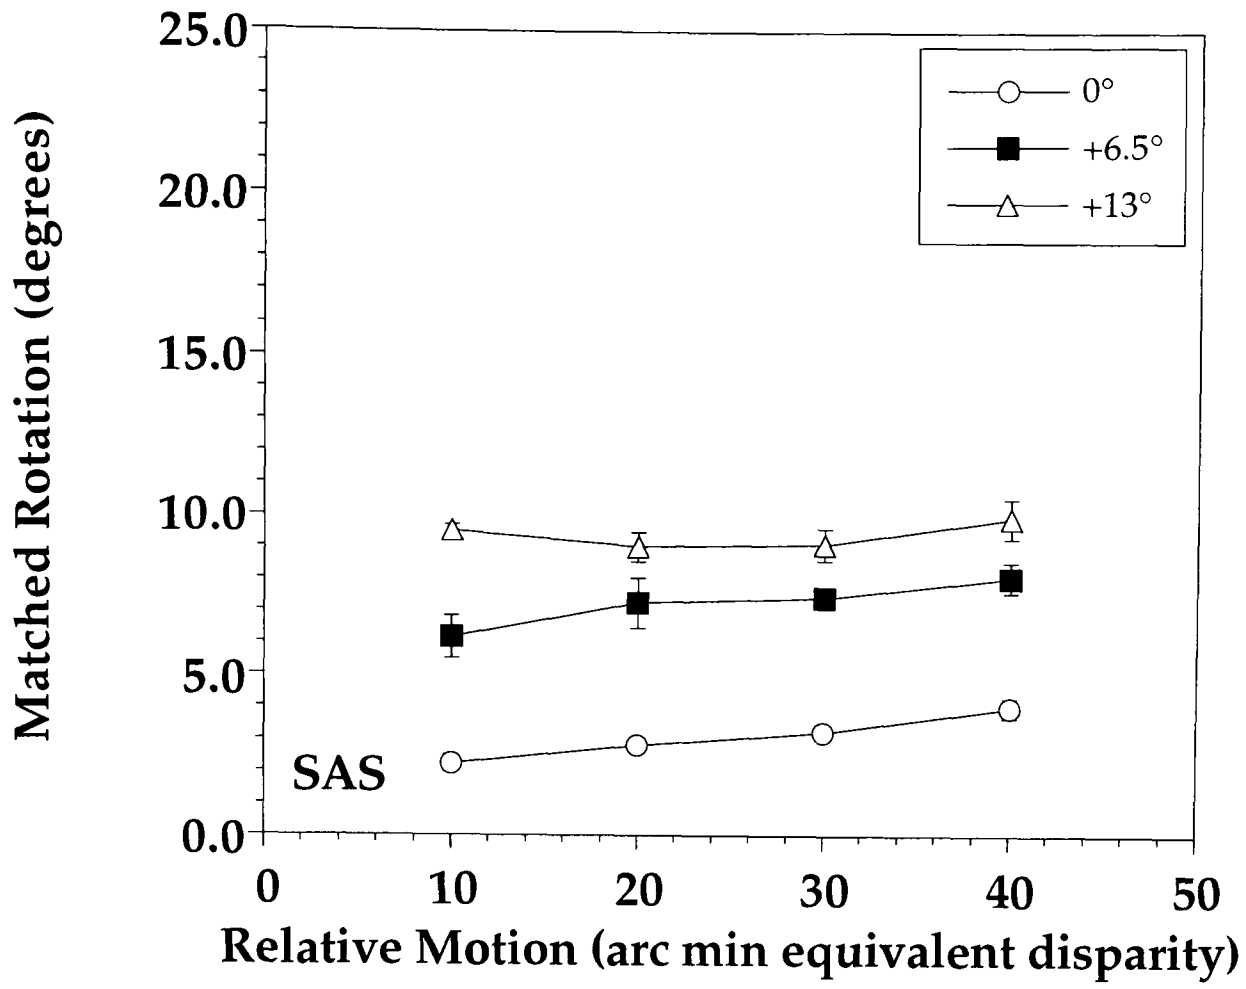


Fig. 6.23: Matched Rotation as a function of relative motion amplitude, for 'right-angled' stimuli for Subject SAS. The vertical perspective cue signalled a rotation angle of either 0°, +6.5°° or +13°.

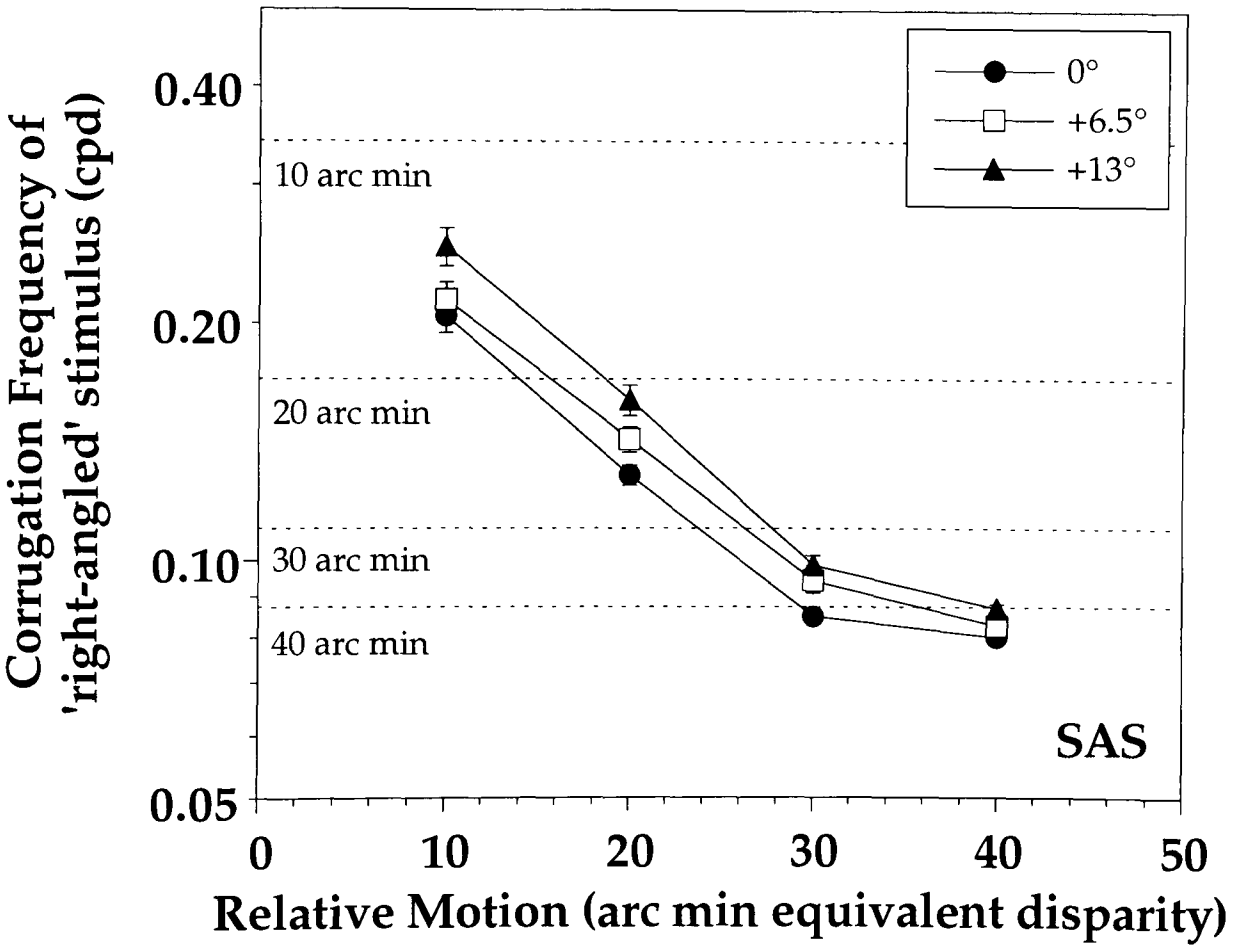


Fig. 6.24: Corrugation frequency at which the stimulus appeared right-angled, as a function of relative motion amplitude for Subject SAS. The rotation angle specified by the vertical perspective cue was either 0°, +6.5° or +13°. The dashed lines show the amplitude at which the surface of each relative motion amplitude is physically right-angled.

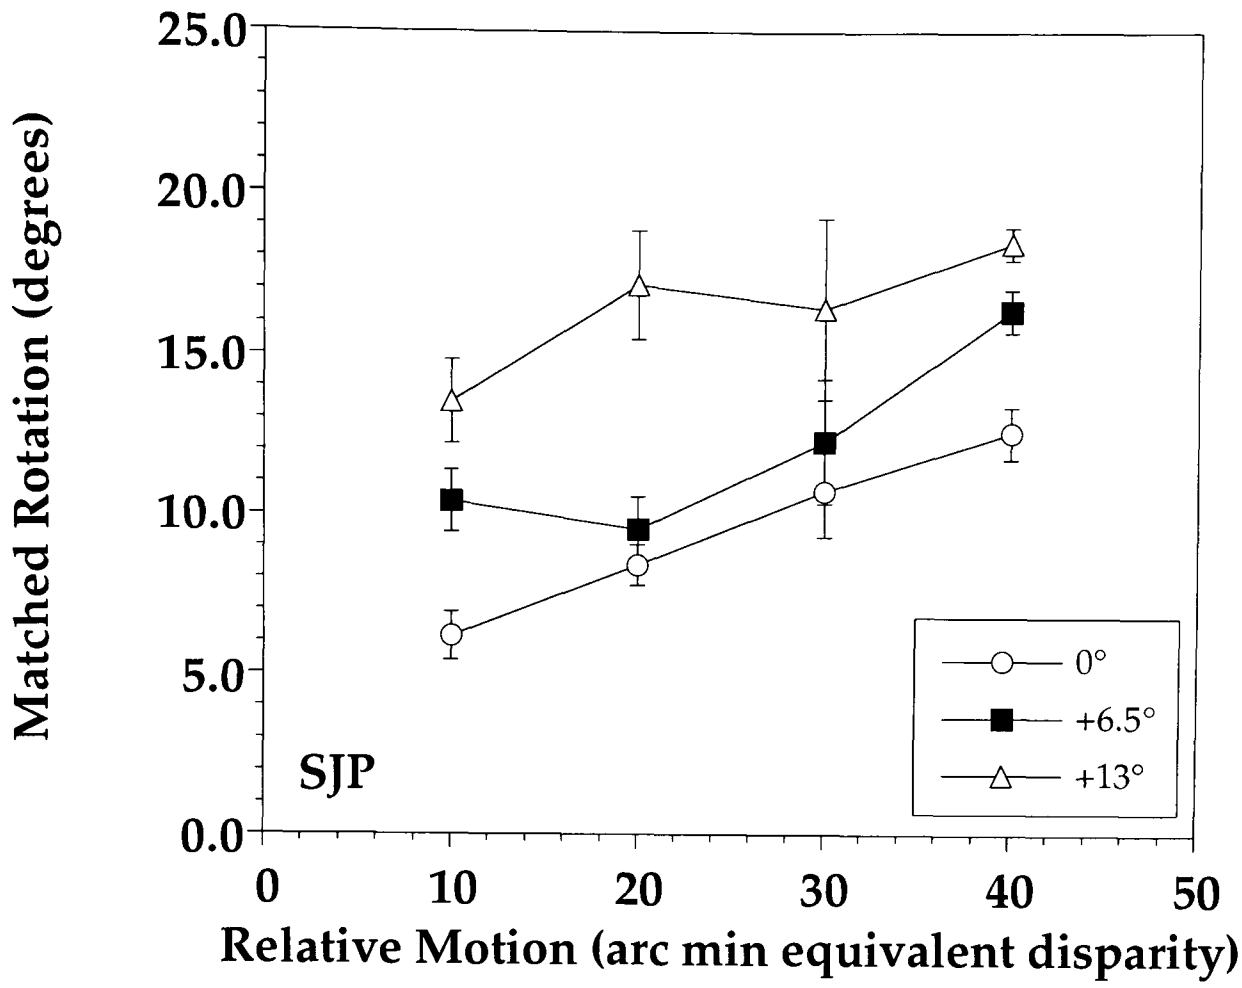


Fig. 6.25: Matched Rotation as a function of relative motion amplitude, for 'right-angled' stimuli for Subject SJP. The vertical perspective cue signalled a rotation angle of either 0°, +6.5° or +13°.

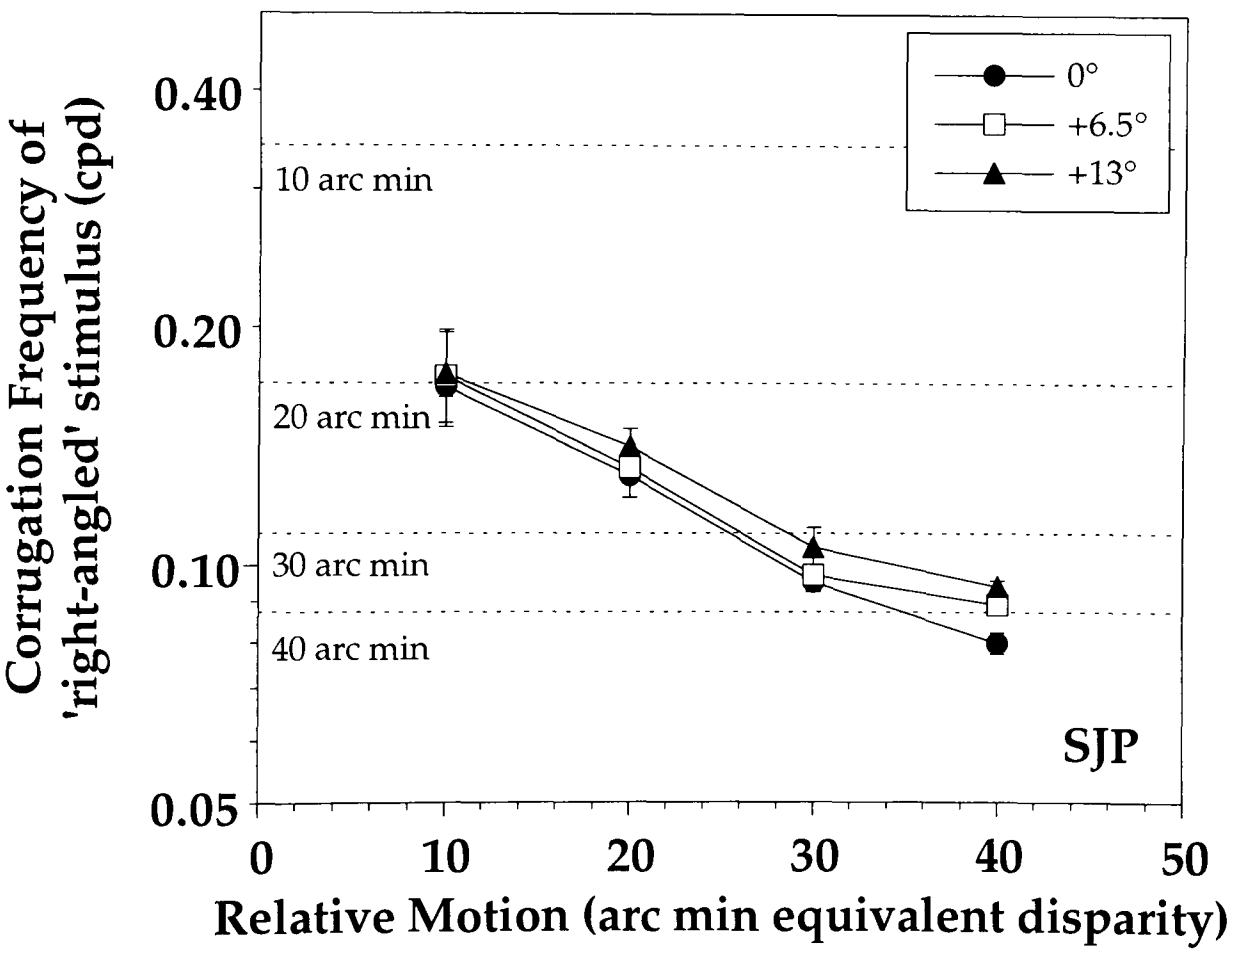


Fig. 6.26: Corrugation frequency at which the stimulus appeared right-angled, as a function of relative motion amplitude for Subject SJP. The rotation angle specified by the vertical perspective cue was either 0°, +6.5° or +13°. The dashed lines show the amplitude at which the surface of each relative motion amplitude is physically right-angled.

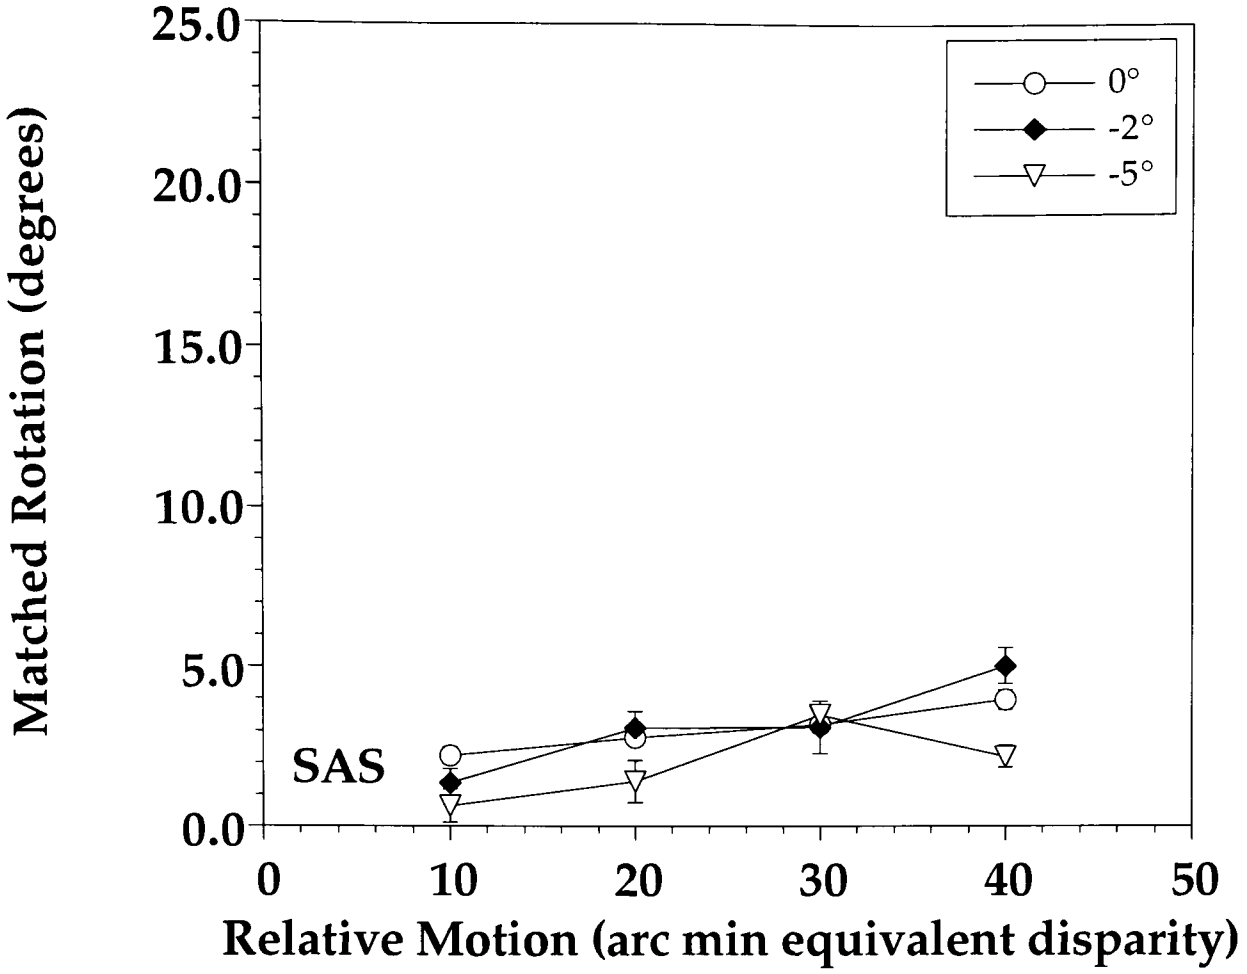


Fig. 6.27: Matched Rotation as a function of relative motion amplitude, for 'right-angled' stimuli for Subject SAS. The vertical perspective cue signalled a rotation angle of either 0°, -2° or -5°.

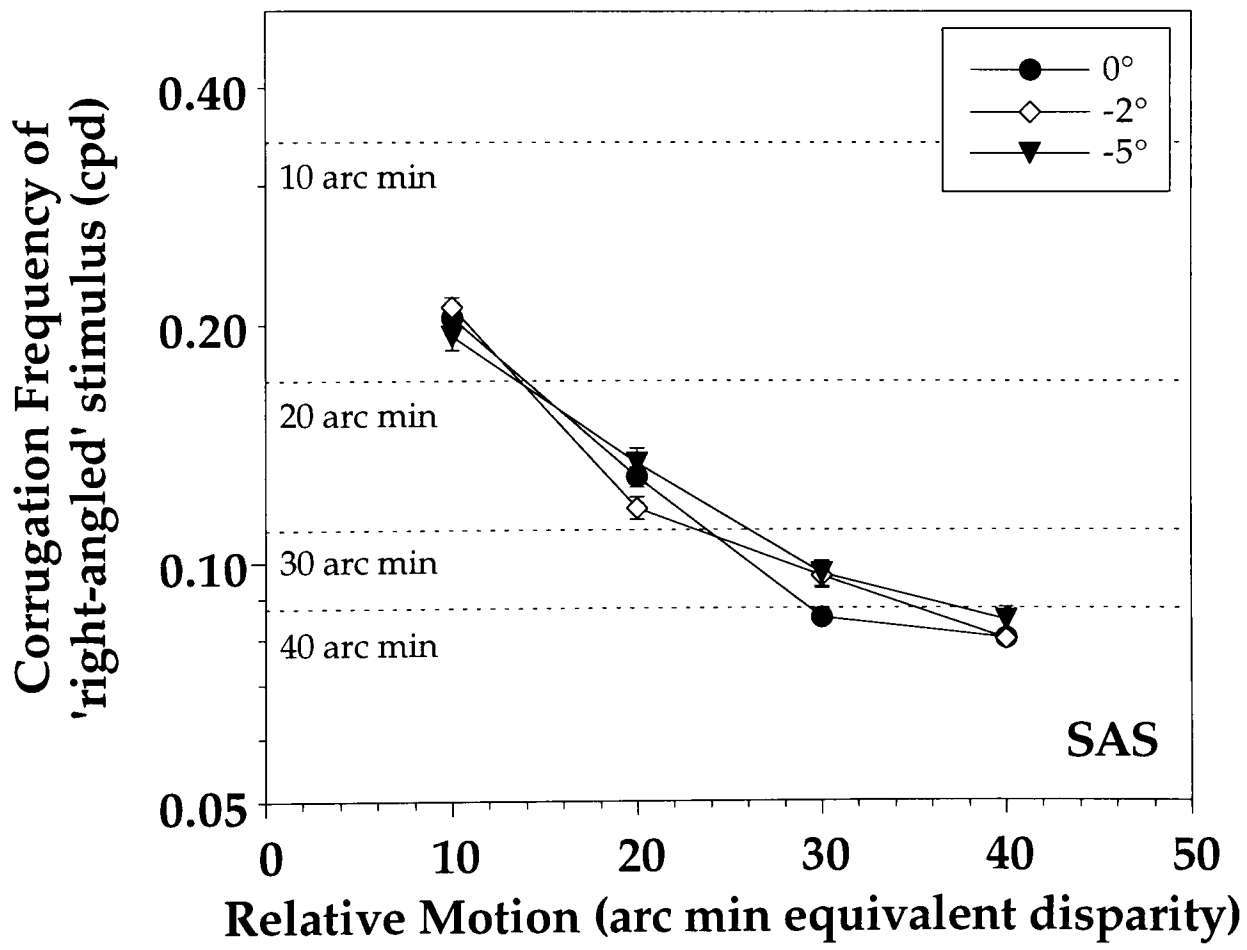


Fig. 6.28: Corrugation frequency at which the stimulus appeared right-angled, as a function of relative motion amplitude for Subject SAS. The rotation angle specified by the vertical perspective cue was either 0°, -2° or -5°. The dashed lines show the amplitude at which the surface of each relative motion amplitude is physically right-angled.

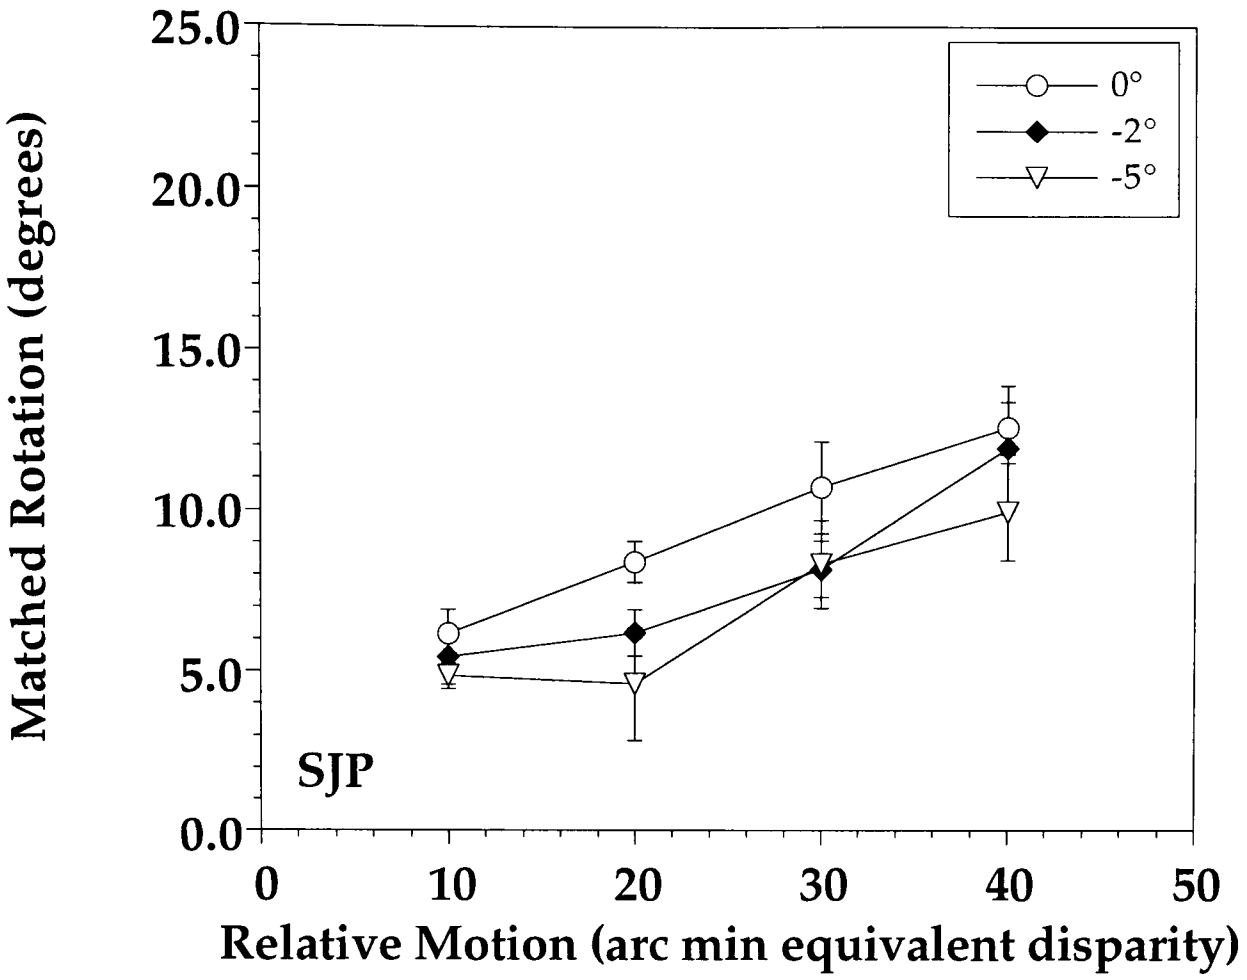


Fig. 6.29: Matched Rotation as a function of relative motion amplitude, for 'right-angled' stimuli for Subject SJP. The vertical perspective cue signalled a rotation angle of either 0°, -2° or -5°.

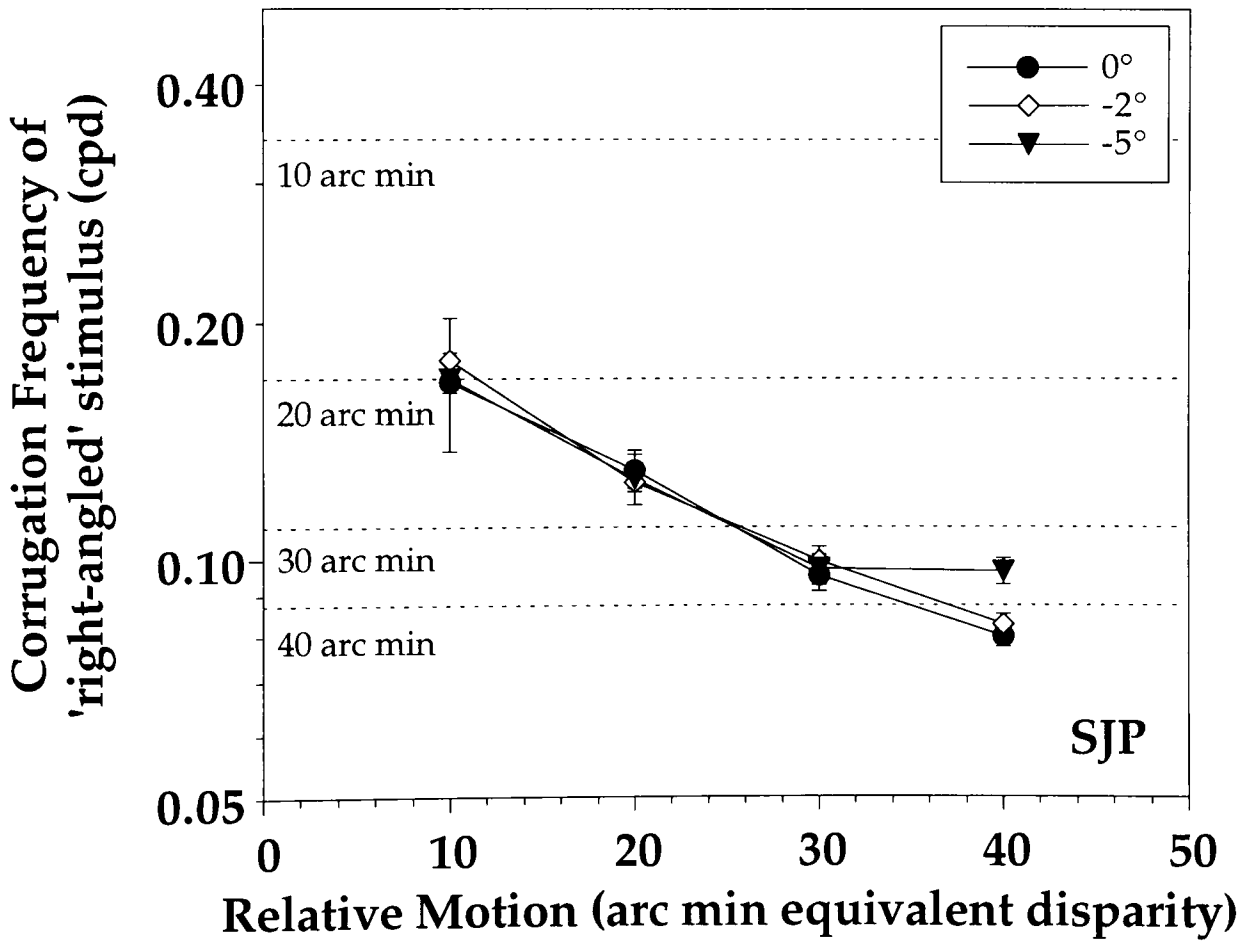


Fig. 6.30: Corrugation frequency at which the stimulus appeared right-angled, as a function of relative motion amplitude for Subject SJP. The rotation angle specified by the vertical perspective cue was either 0°, -2° or -5°. The dashed lines show the amplitude at which the surface of each relative motion amplitude is physically right-angled.

6.7 Experiment 20: The Effect of Dynamic Vertical Perspective on Matched Depth and Matched Rotation as a function of Viewing Distance

6.7.1 Introduction

The aim of the current experiment was to examine the effectiveness of the vertical perspective cue on the perceived magnitude of rotation, as a function of viewing distance. Experiment 15 showed that the amount of perceived rotation of proximally identical surfaces increases with increasing viewing distance. A possible reason for this, suggested in the discussion following that experiment, is that as viewing distance increases, the visual system becomes decreasingly constrained by the vertical perspective information which specifies that no rotation is occurring. It may thus be more willing to attribute the relative motion to rotation.

The question remains as to why the visual system should take less account of vertical perspective information with increasing viewing distance. For a stimulus of constant angular size which rotates through a constant angle, the absolute size of the changes in vertical perspective are independent of viewing distance. There would thus seem to be no *a priori* reason for an increase in the perceived rotation magnitude of the surface with increasing viewing distance. In the real world, however, the vertical perspective changes accompanying the rotation of an object through a constant angle *decrease* with increasing viewing distance, due to the decreasing proximal size of the distally constant stimulus. It might be the case that the visual system becomes decreasingly reliant upon vertical perspective information (with increasing viewing distance), due to its decreasing absolute size and consequently its increasingly difficult extraction. The increasing propensity to perceive rotation in proximally identical stimuli with increasing viewing distance might simply reflect this putative bias of the visual system to take less account of vertical perspective information at greater distances.

The current experiment examined the perceived rotation of surfaces in which the magnitude of rotation simulated by the vertical perspective cue took one of two values at each viewing distance. The simulated rotation magnitudes remained constant with viewing distance. If vertical

perspective information plays a decreasing role in the perceived magnitude of rotation with increasing viewing distance, then it would be expected that the difference between the perceived rotation of the two surfaces would decrease with increasing viewing distance.

Concomitant variations in perceived depth magnitude were also measured.

6.7.2 Stimuli

There were two sets of stimuli. The stimuli within each set were proximally identical in terms of corrugation frequency, angle subtended by each display element, display size and amplitude of horizontal relative motion. The first set consisted of stimuli which subtended 5° and which contained relative motion amplitudes of 2 - 10 arc min equivalent disparity. These stimuli were viewed from distances of 57, 114 and 228 cm. The second set consisted of stimuli which subtended 10° and which contained relative motion amplitudes of 4 - 20 arc min equivalent disparity. These stimuli were viewed from distances of 57 and 114 cm.

The magnitude of rotation simulated by the vertical perspective cue was either 0° or ±15°. All stimuli had a corrugation frequency of 0.4 cpd.

6.7.3 Procedure

Trials were blocked in groups of ten, corresponding to two repetitions of each of the relative motion amplitudes. Each experimental session consisted of ten blocks [(three viewing distances x two vertical perspective levels, for the size 5° stimuli) + (two viewing distances x two vertical perspective levels, for the size 10° stimuli)]. Subjects completed three experimental sessions. The order of blocks within each session was randomised.

The rotation and depth matching procedures were identical to those of Experiment 16.

6.7.4 Results

Figs. 6.31 - 6.48 show Matched Rotation and Matched Depth as a function of relative motion amplitude for size 5° stimuli at viewing distances of 57, 114 and 228 cm. Open symbols represent matches for stimuli in which the simulated rotation magnitude was 0°. Closed symbols represent matches for stimuli in which the simulated rotation magnitude was 15°. The right-hand ordinate of each depth matching graph shows simulated depth in mm.

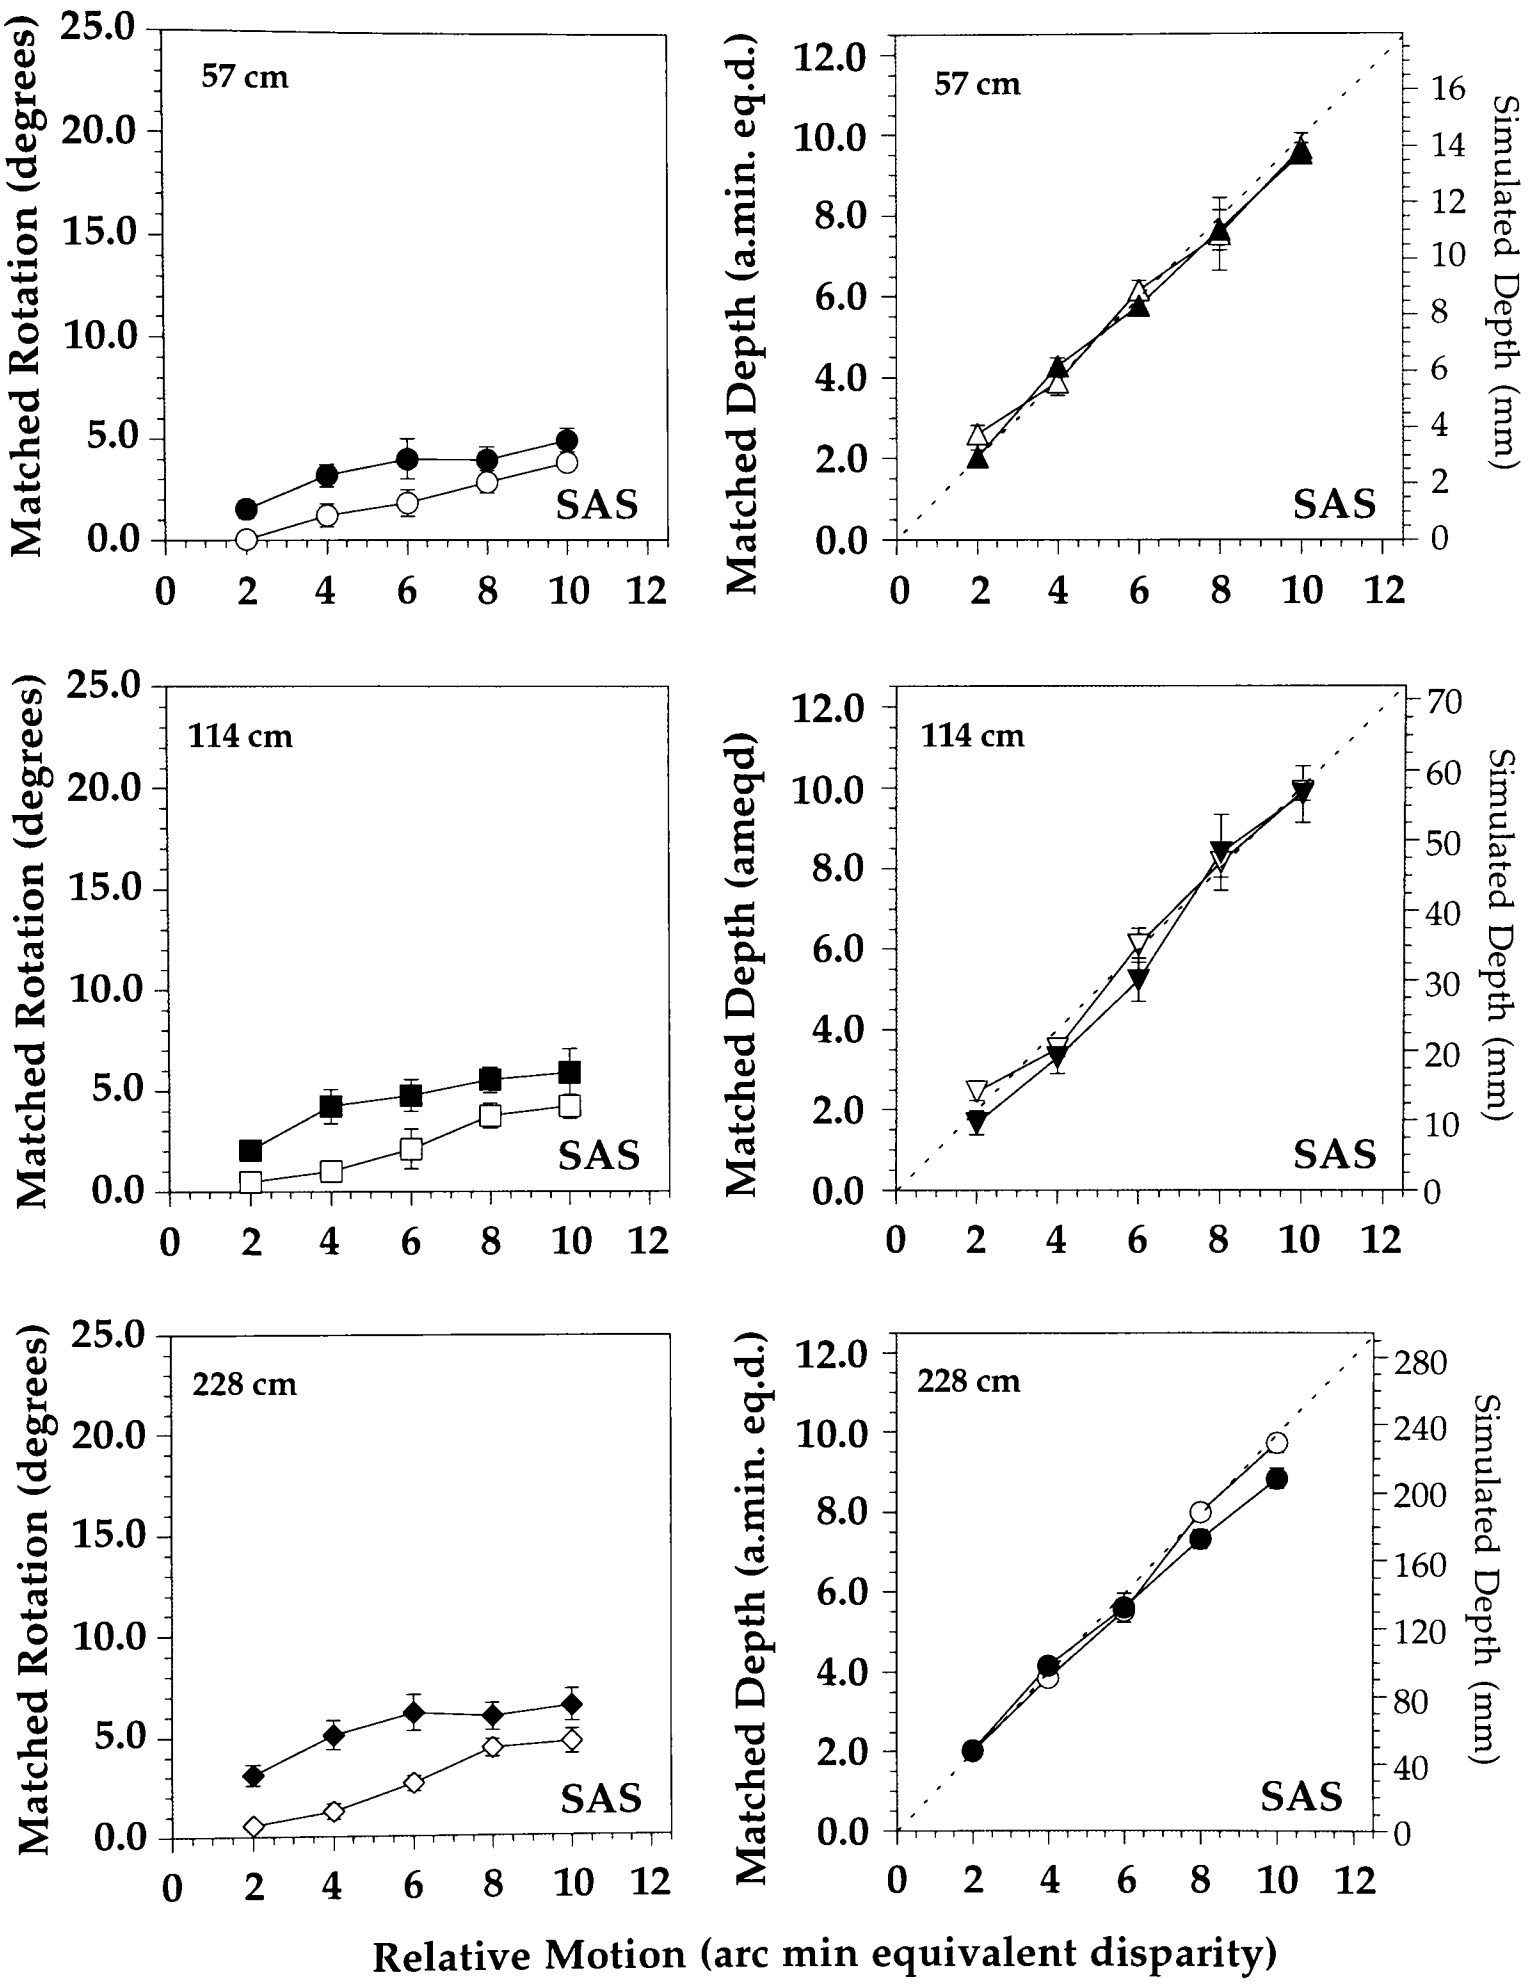


Fig. 6.31 - 6.36: Matched Rotation (left-hand panels) and Matched Depth (right-hand panels) as a function of Relative Motion for surfaces of size 5° at viewing distances of 57, 114 and 228 cm for Subject SAS. Open symbols represent the condition in which vertical perspective simulated a rotation angle of 0°. Closed symbols represent the condition in which vertical perspective simulated a rotation angle of 15°.

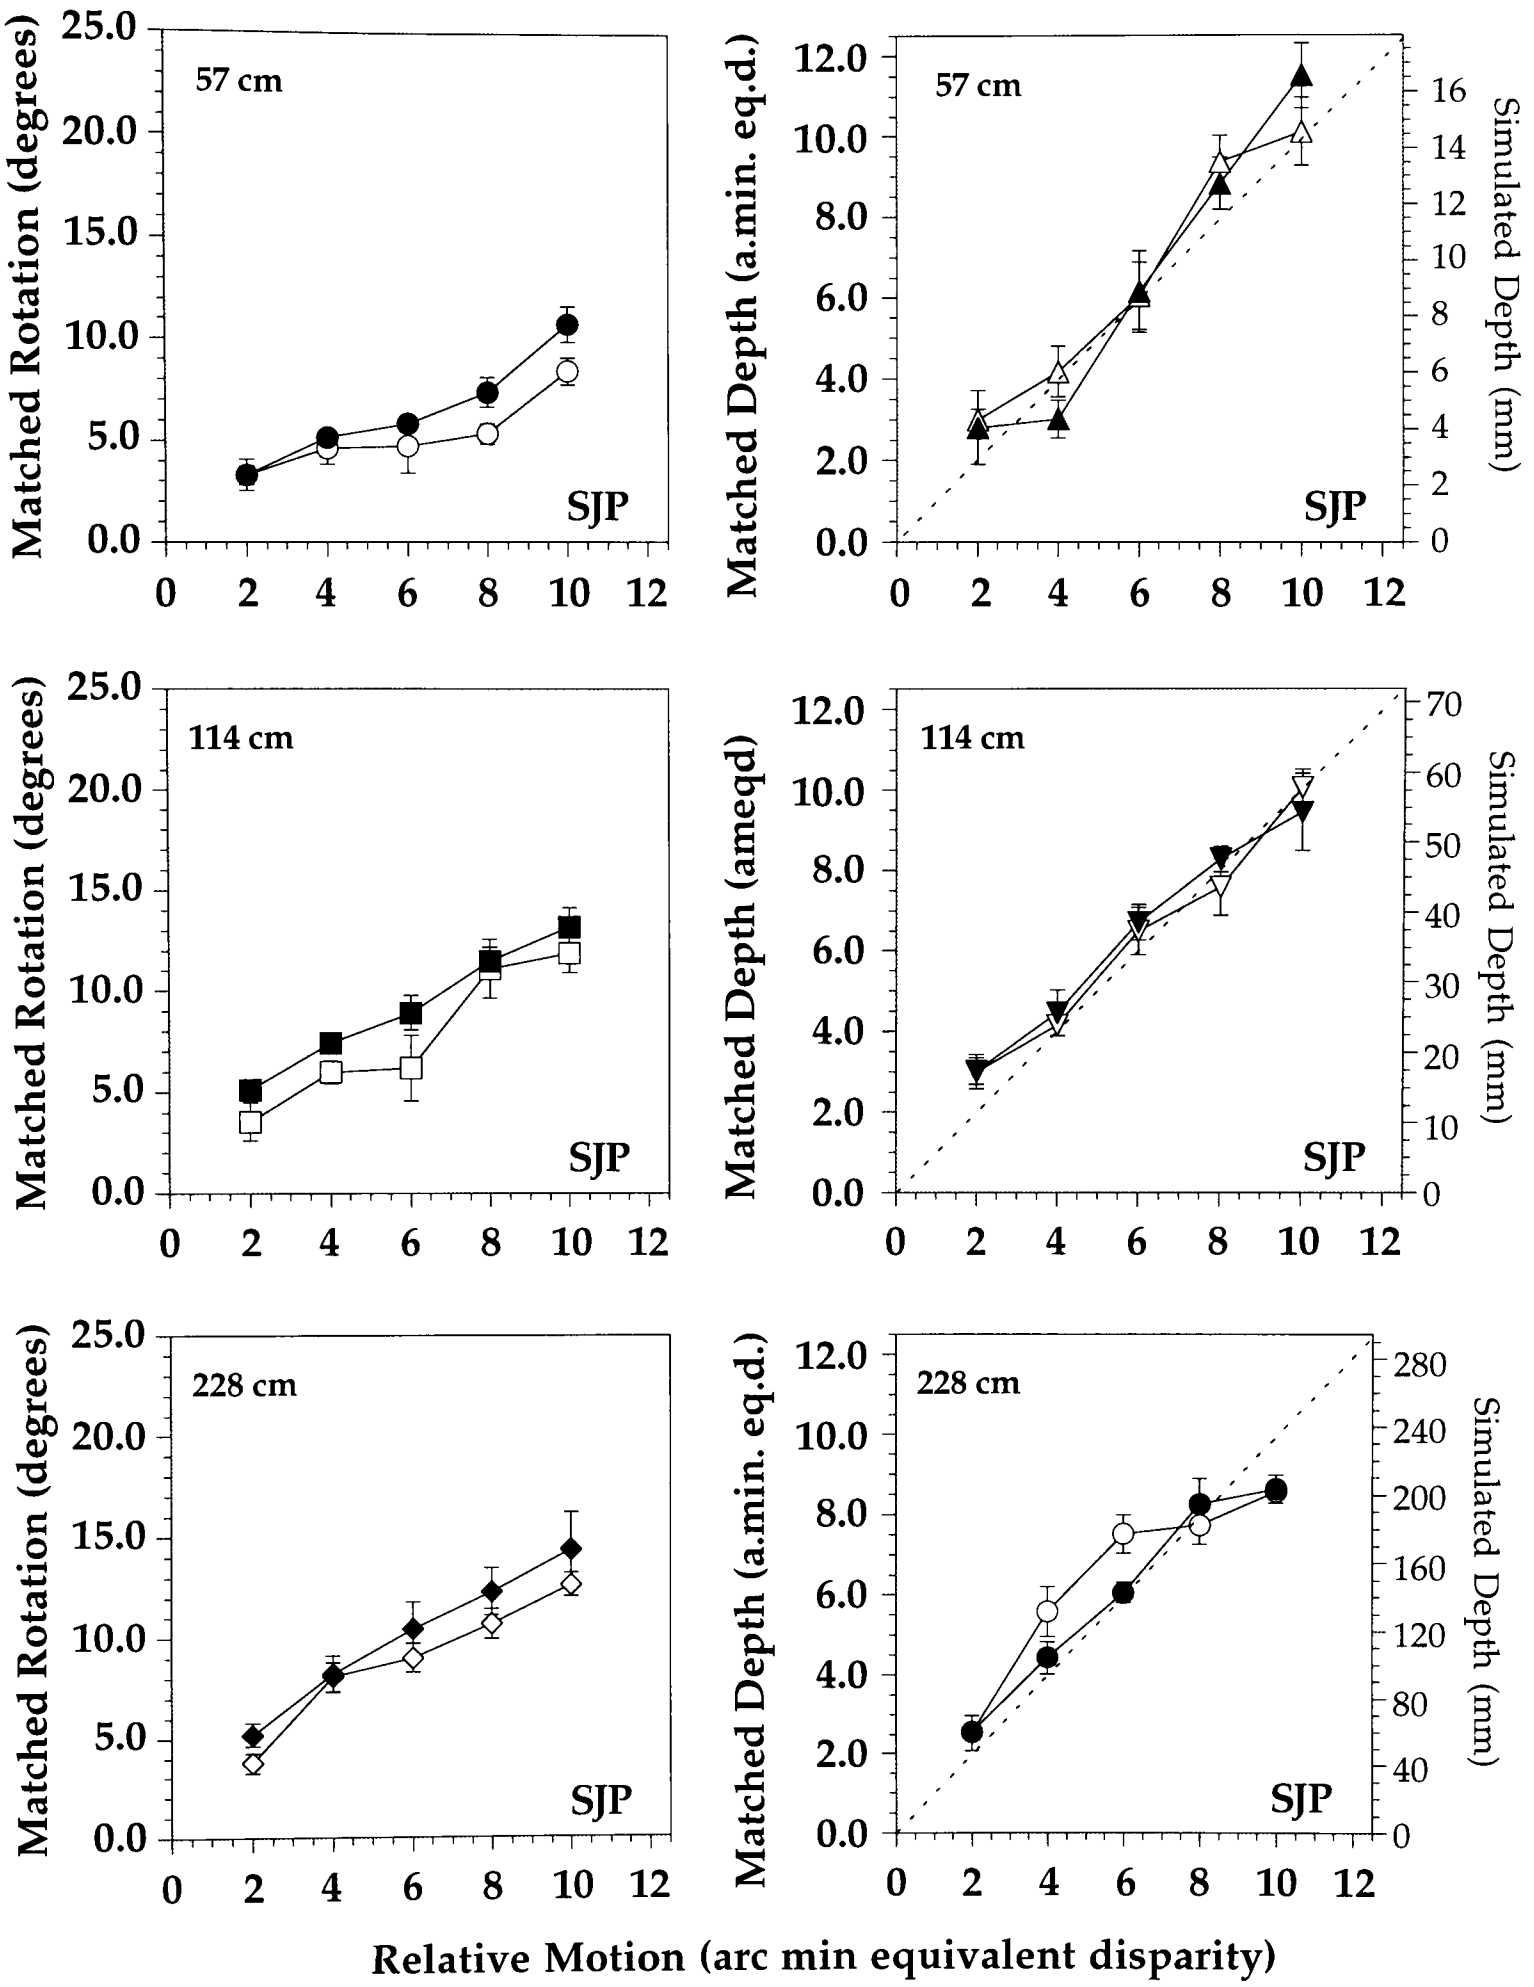


Fig. 6.37 - 6.42: Matched Rotation (left-hand panels) and Matched Depth (right-hand panels) as a function of Relative Motion for surfaces of size 5° at viewing distances of 57, 114 and 228 cm for Subject SJP. Open symbols represent the condition in which vertical perspective simulated a rotation angle of 0°. Closed symbols represent the condition in which vertical perspective simulated a rotation angle of 15°.

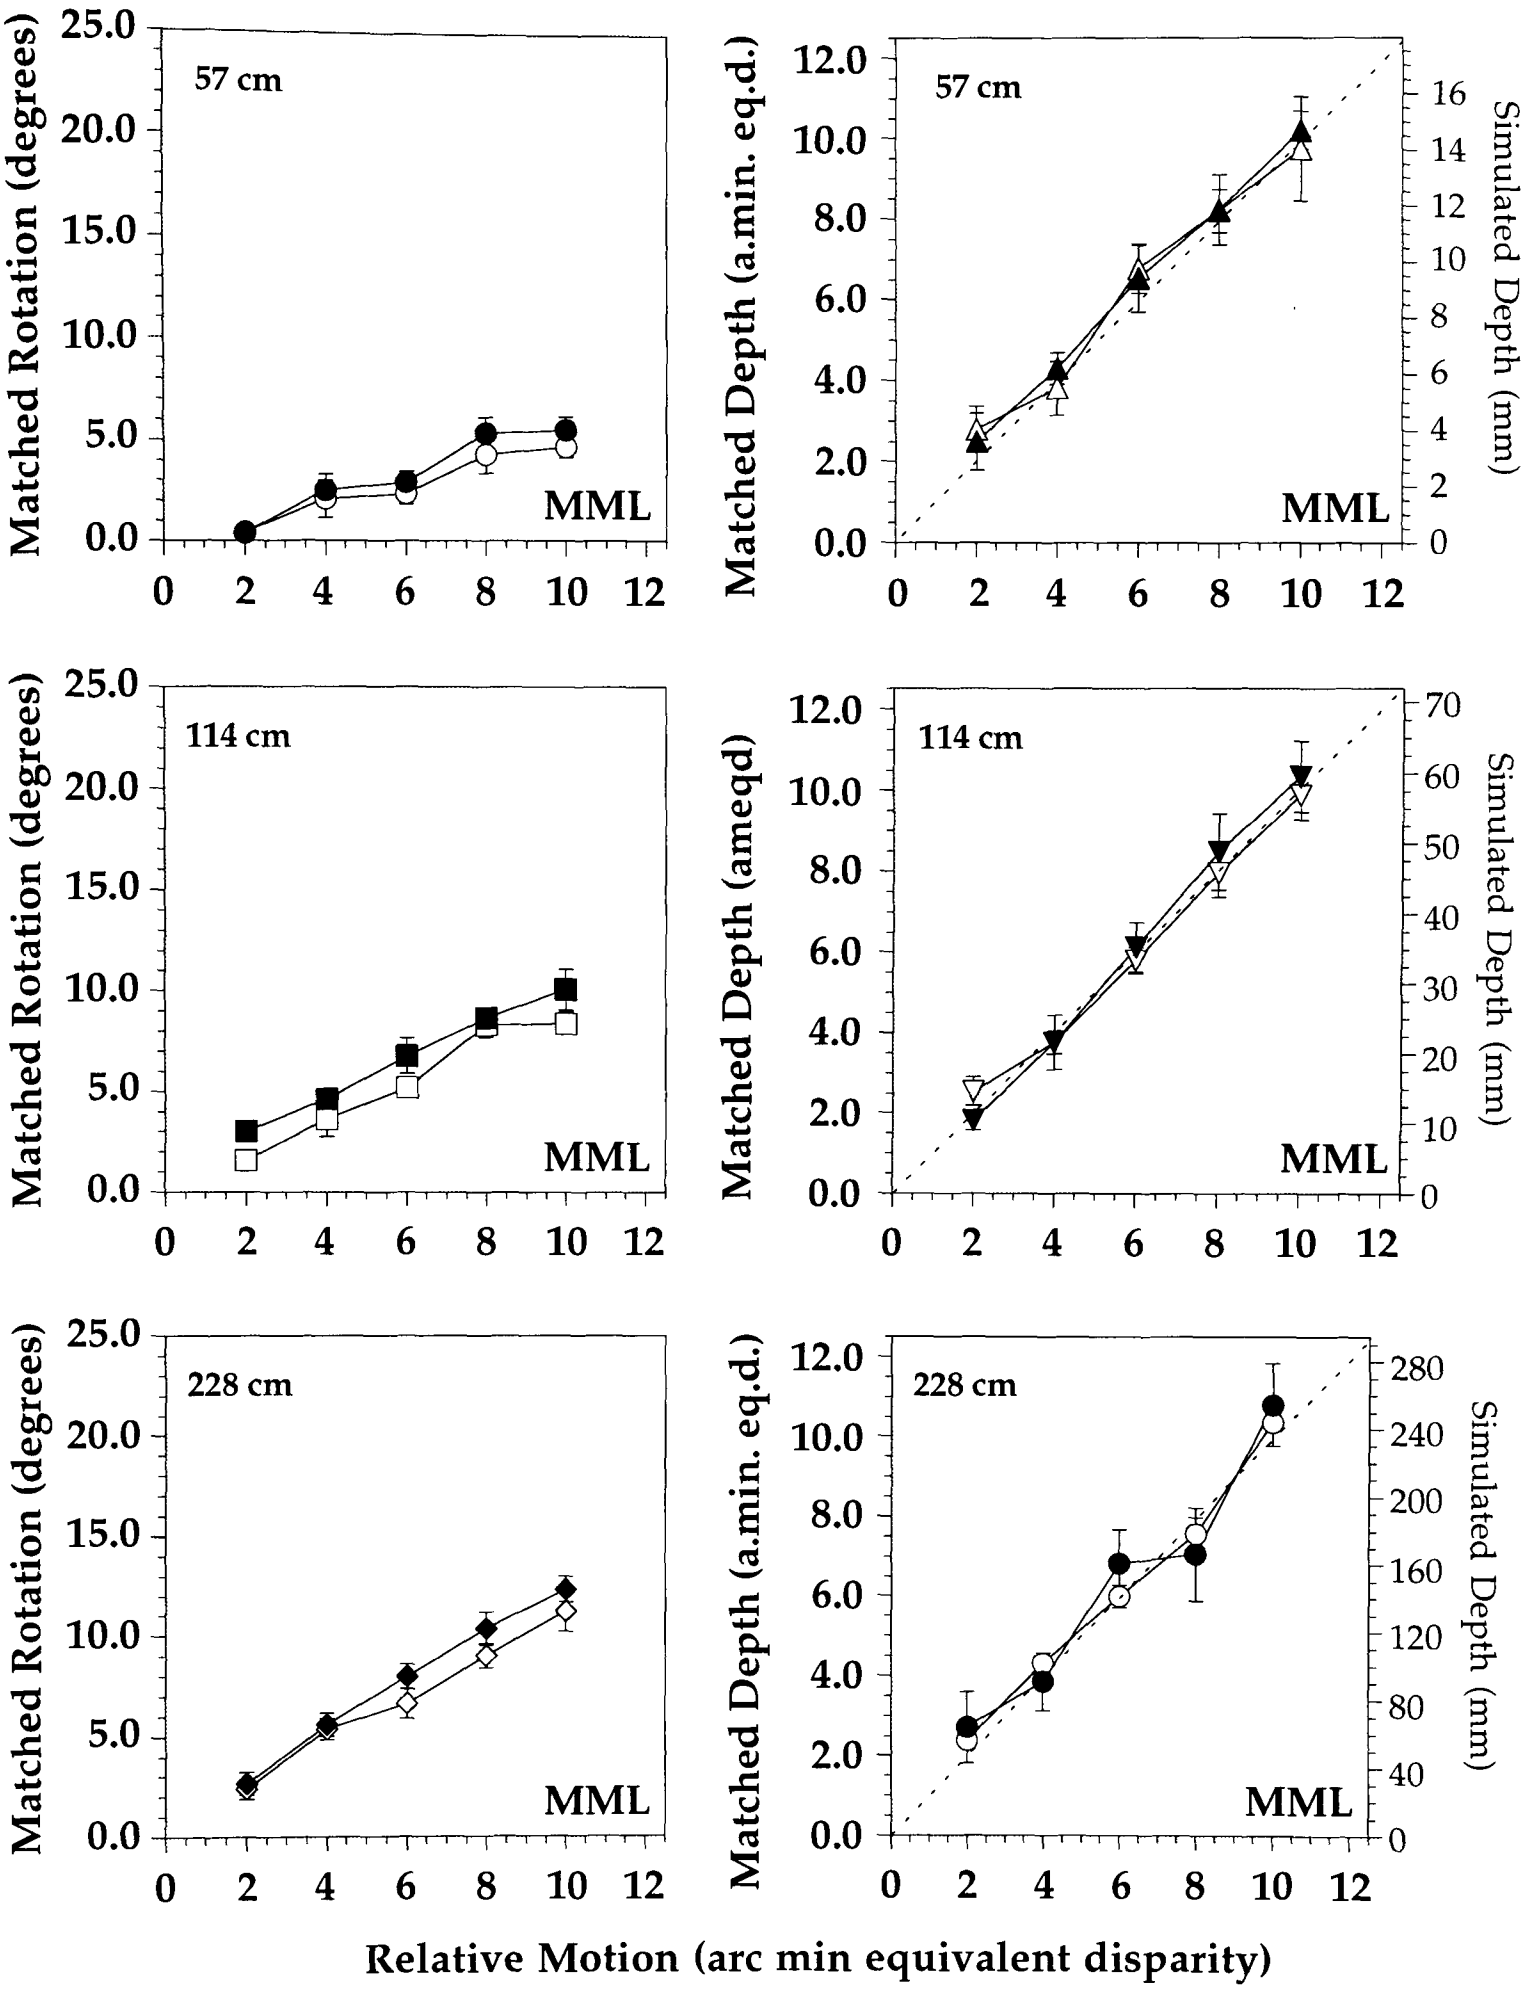


Fig. 6.43 - 6.48: Matched Rotation (left-hand panels) and Matched Depth (right-hand panels) as a function of Relative Motion for surfaces of size 5° at viewing distances of 57, 114 and 228 cm for Subject MML. Open symbols represent the condition in which vertical perspective simulated a rotation angle of 0°. Closed symbols represent the condition in which vertical perspective simulated a rotation angle of 15°.

Figs. 6.49 - 6.60 show Matched Rotation and Matched Depth as a function of relative motion amplitude for size 10° stimuli at viewing distances of 57 and 114 cm. Open symbols represent matches for stimuli in which the simulated rotation magnitude was 0°. Closed symbols represent matches for stimuli in which the simulated rotation magnitude was 15°. Again, the right-hand ordinate of each depth matching graph shows simulated depth in mm.

Considering first the rotation matches for the stimuli of size 5° which contained relative motion amplitudes of 2 - 10 arc min equivalent disparity, it can be seen that the effect of the vertical perspective cue was similar at each viewing distance. Although the absolute values of the matches in each vertical perspective condition increased with viewing distance, the relative positions of the functions remained approximately constant. The depth matches were not significantly affected by the vertical perspective cue.

The rotation matches for the size 10° stimuli which contained relative motion amplitudes of 4 - 20 arc min equivalent disparity were, again, slightly greater at the further viewing distance. The effect of vertical perspective was generally similar at each viewing distance for Subjects SJP and MML, although Subject SAS appeared to be affected by the perspective cue slightly more at the greater viewing distance. The depth matches were slightly less in the 15° conditions than in the 0° conditions. As found in Experiments 16 - 19, this difference was very small.

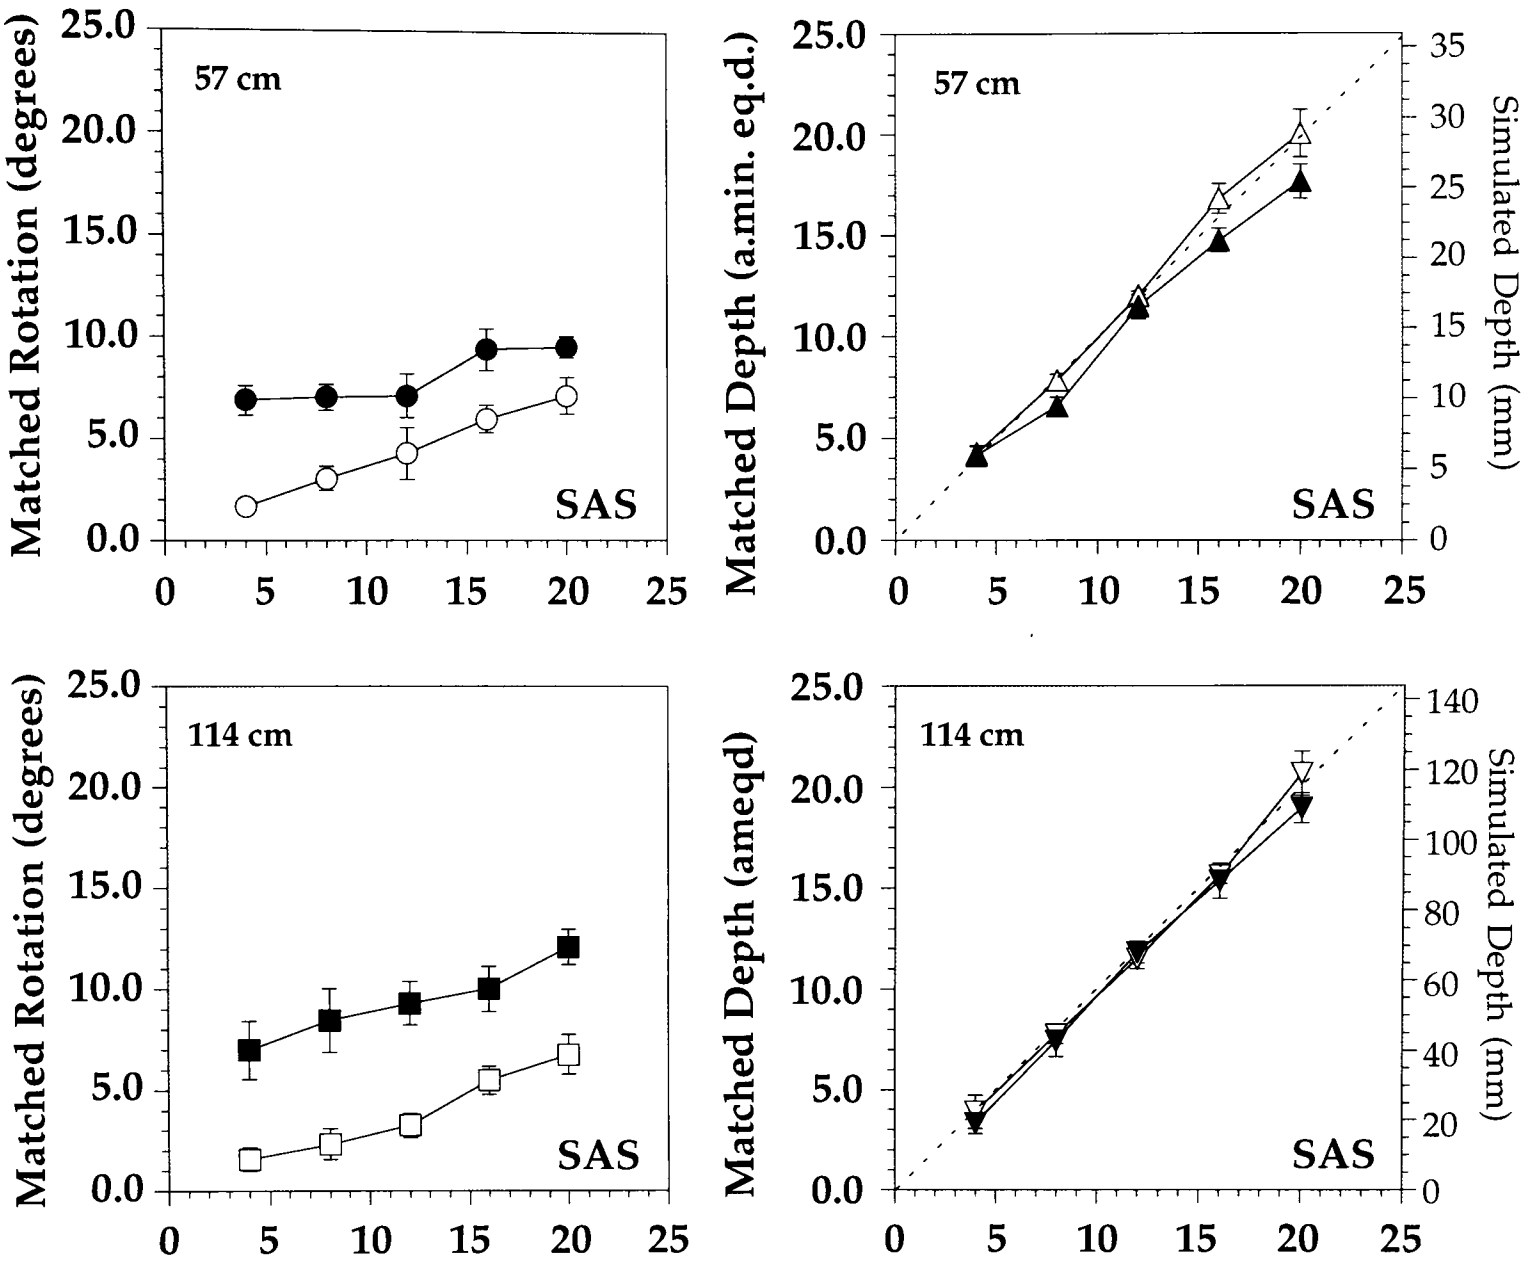


Fig. 6.49 - 6.52: Matched Rotation (left-hand panels) and Matched Depth (right-hand panels) as a function of Relative Motion for surfaces of size 10° at viewing distances of 57 and 114 cm for Subject SAS. Open symbols represent the condition in which vertical perspective simulated a rotation angle of 0°. Closed symbols represent the condition in which vertical perspective simulated a rotation angle of 15°.

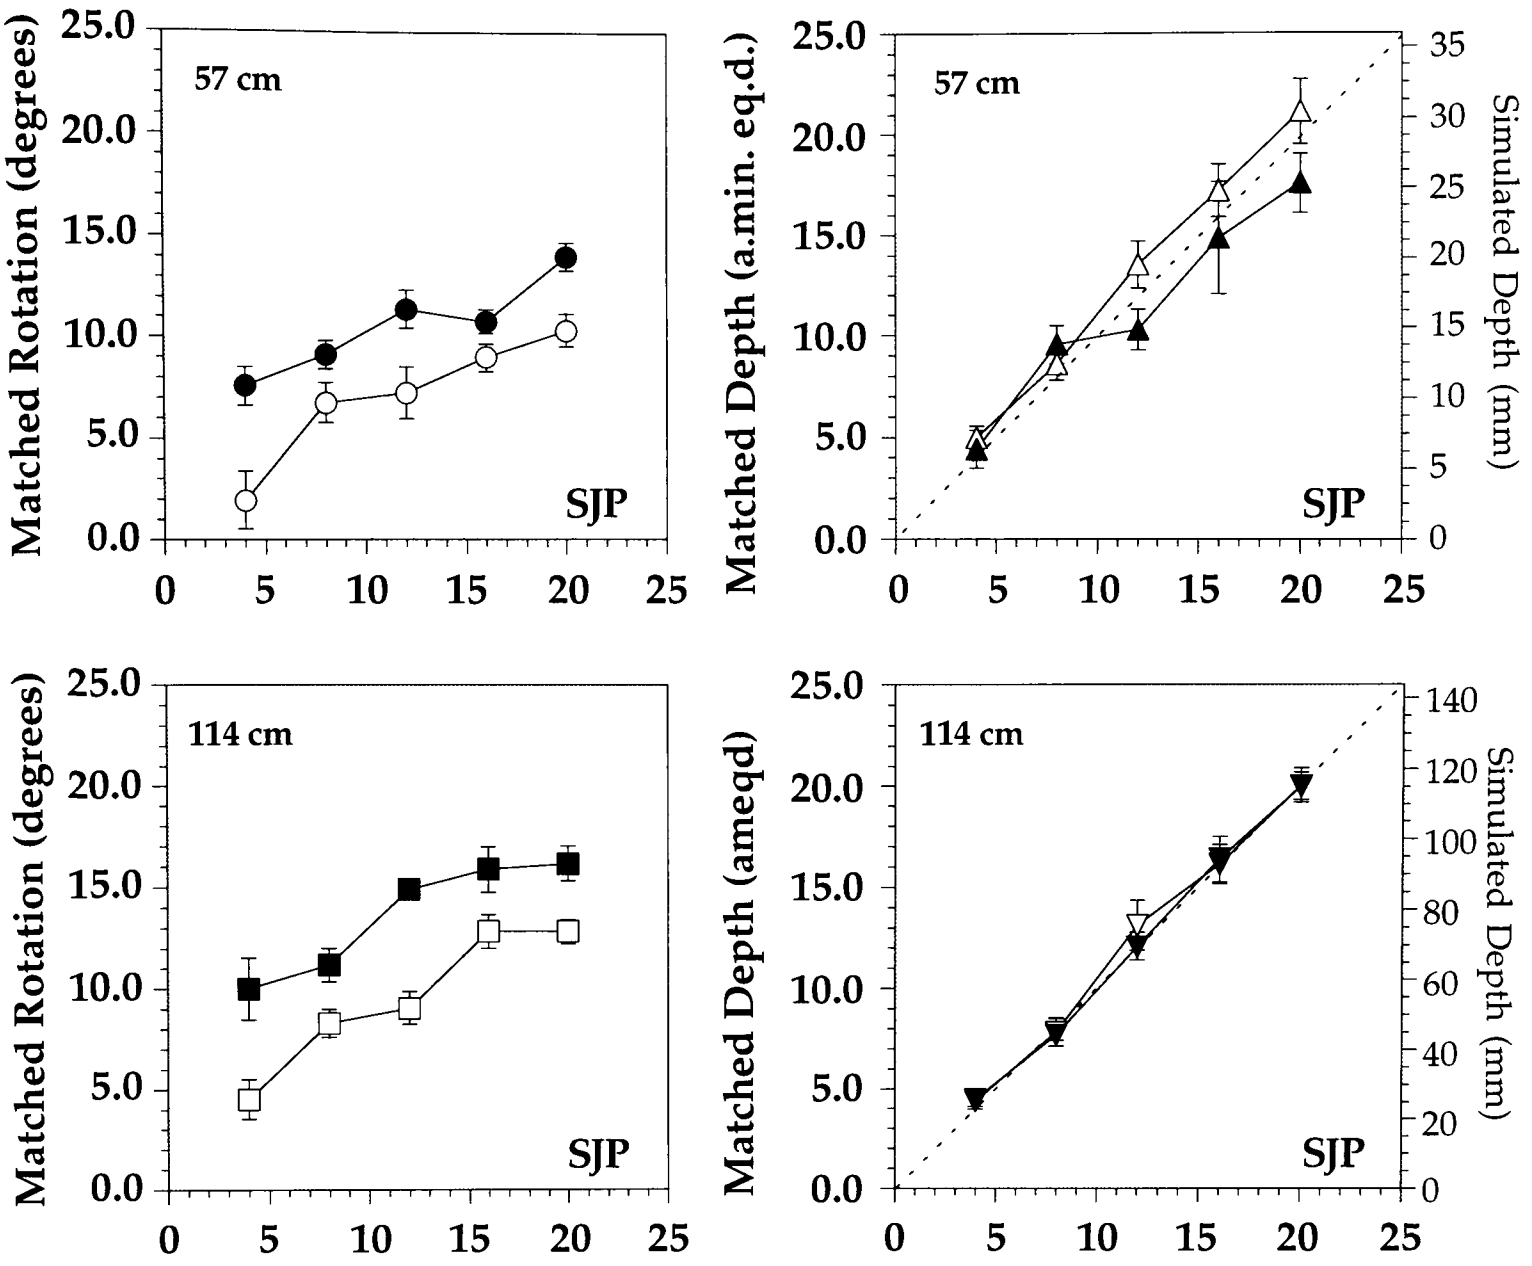


Fig. 6.53 - 6.56: Matched Rotation (left-hand panels) and Matched Depth (right-hand panels) as a function of Relative Motion for surfaces of size 10° at viewing distances of 57 and 114 cm for Subject SJP. Open symbols represent the condition in which vertical perspective simulated a rotation angle of 0°. Closed symbols represent the condition in which vertical perspective simulated a rotation angle of 15°.

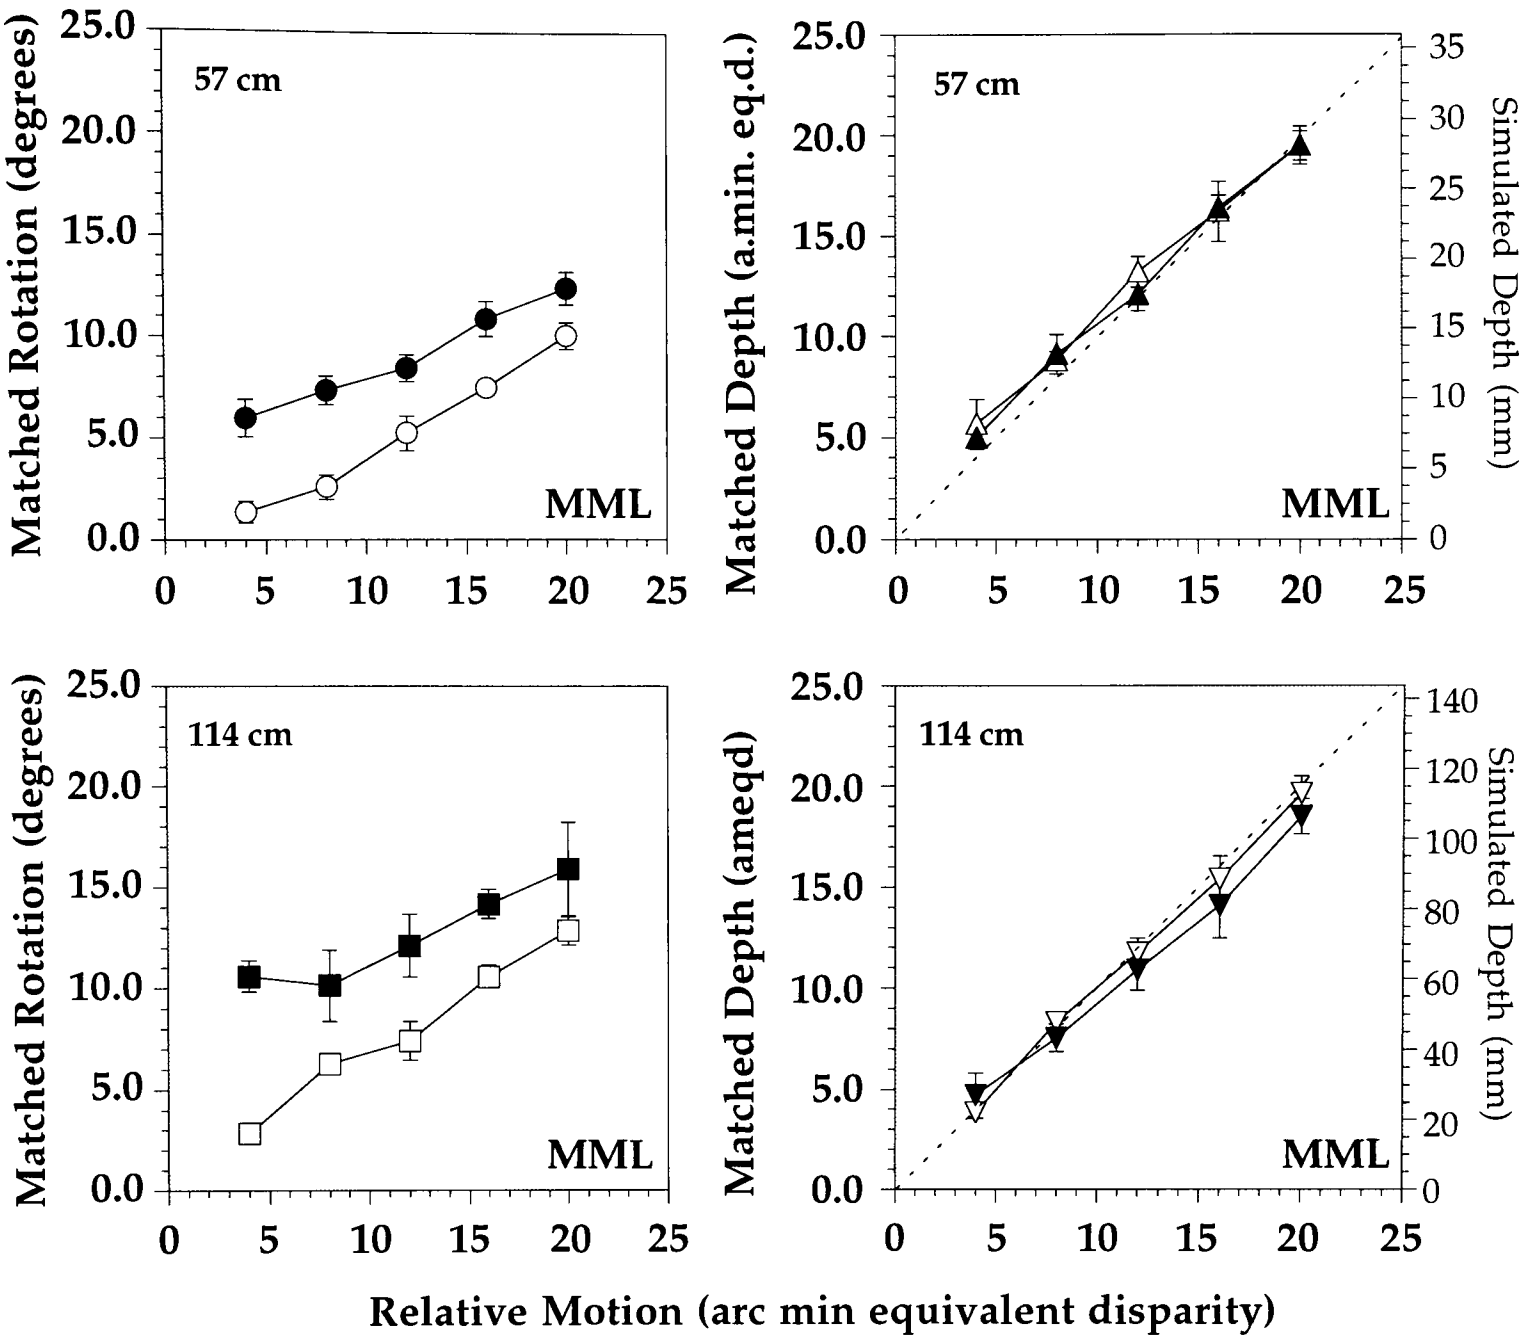


Fig. 6.57 - 6.60: Matched Rotation (left-hand panels) and Matched Depth (right-hand panels) as a function of Relative Motion for surfaces of size 10° at viewing distances of 57 and 114 cm for Subject MML. Open symbols represent the condition in which vertical perspective simulated a rotation angle of 0°. Closed symbols represent the condition in which vertical perspective simulated a rotation angle of 15°.

6.7.5 Discussion

This experiment did not provide evidence to support the hypothesis that vertical perspective information is increasingly ignored as viewing distance increases.

A factor which could have played a role in the observed pattern of results is the effect of observer translation upon the vertical perspective changes in an object's retinal image, as a function of viewing distance. Consider a stationary surface in the frontal plane — the face of the oscilloscope, for instance — viewed by an observer translating laterally

through a given distance. The change in obliqueness of the surface with respect to the eye which occurs with the observer's translation does *not* remain constant with variations in viewing distance (see Section 6.1). For each doubling of distance, the changes in surface obliqueness which accompany the observer's translation are halved. If the visual system was presented with a stimulus containing vertical perspective changes which remained constant over distance, given that — for a given extent of observer translation — the vertical perspective changes which arise *due to the observer's own translation* decrease with increasing viewing distance, it might be the case that the visual system would be increasingly likely, with increasing viewing distance, to attribute this perspective change to object rotation. This would predict the results of Experiment 15. If it is assumed that the increased propensity to attribute vertical perspective changes to object rotation rather than observer translation does not vary with the simulated angle of rotation, then it would also predict the results of the present experiment. Although predicting that for a given simulated surface rotation the perceived rotation magnitude would increase with increasing viewing distance, it predicts that the *difference* in the perceived magnitudes of rotation at the two different levels of vertical perspective would remain constant over viewing distance.

A notable feature of these data is the *prima facie* near-perfect depth scaling which occurred with variations in viewing distance. This is in apparent contrast to the results of Experiment 15, which showed that depth was increasingly under-estimated as viewing distance increased. It will be noted, however, that the techniques used to measure perceived depth in the two cases were dissimilar. Electronic callipers were used in the former experiment, while an electronic match stimulus was employed in the current experiment. The reason for the introduction of the latter technique was discussed in Section 6.2.3.

As noted by Glennerster et al. (1996), different depth matching techniques may give rise to dissimilar results because of differences between the visual processing of the observer in the two cases. The authors found that stereoscopic depth constancy was 75% when observers were required to set the depth of a hemi-cylinder to equal its half-height, or a dihedral angle to 90°, but was 100% when the task was to match the depth of sinusoidal corrugations presented at two distances. The authors

suggested that the former task requires an estimate of the viewing distance, in order to obtain information about the Euclidean structure of the stimulus. It was argued that the latter task might be performed using a simple 'direct' strategy, which did not require the construction and comparison of metric representations of each surface. This study highlights the fact that different results may be obtained when different matching tasks are employed, and it is possible that the discrepancy between results obtained in the current experiment and those of Experiment 15 are at least partly attributable to the differing tasks used. It is possible that depth matching using the callipers involves an internal metric representation of depth, while the similarity of the electronic match and test stimuli of Experiment 20 may have allowed a more 'direct' strategy, perhaps along the lines of that suggested by Glennerster et al. (1996), to be used.

It will also be recalled that the electronic match stimulus was always presented at the same viewing distance as the test stimulus; depth scaling was not directly measured in the current experiment for this very reason. Consequently, it is unsurprising that the simulated and matched depths were similar despite the increase in viewing distance. The primary purpose of the current experiment was not to measure depth constancy, but rather was to compare the *relative* amounts of matched depth between stimuli containing differing amounts of vertical perspective changes.

6.8 Experiment 21: The Effect of Magnitude of Dynamic Vertical Perspective and Stimulus Size on Perceived Depth Order in Externally-Produced Motion Parallax Surfaces

6.8.1 Introduction

The depth order within structures under parallel projection is geometrically ambiguous. This is because the projected motions of the various parts of the structure are independent of their relative distances from the projection surface. The 3-D structure is therefore defined only up to a reflection in the image plane, and its projection is an 'object-centred'

representation because information about relative depth (with respect to the observer) is absent. Observers viewing such stimuli perceive one depth order approximately as frequently as they perceive the other (Rogers and Rogers, 1992).

Early reports of motion as an indicator of depth often involved situations in which the structure of an object was adequately perceived but its direction of rotation and accompanying depth order were not. Examples of such misperceptions include the windmill, fan and rotating trapezoid illusions. Sinsteden (1860) (referenced in Boring, 1942) reported that as he viewed a windmill silhouetted against the evening sky, the direction of rotation of the windmill's arms appeared to reverse. This occurred repeatedly, and it was impossible to tell whether he was viewing the building from its front or its back. Indeed, by imagining on which side of the windmill he was standing, Sinsteden could control the perceived direction of rotation. Miles (1929) reports a practical problem caused by this illusion. A number of purchasers of windmills believed that the arms were actually reversing their direction, and complained to the suppliers that the equipment was faulty.

A similar effect from a two-bladed fan was noted by Kenyon (1898). He reported that the perceived direction of rotation was sometimes opposite to the actual direction. The effect generally occurred when the fan was viewed from approximately 30 feet and the blades were moving at a moderate speed. Miles (1931), recognising that the effect paralleled the 'windmill illusion', investigated it more thoroughly using a shadow projection technique. A two-bladed fan was placed between the light source and a milk glass window. The observer viewed the shadow of the fan from the other side of the window. All but two of the 27 naïve observers reported seeing rotary motion in each direction during the trial.

These instances of illusory directions of rotation are probably largely due to the paucity of cues such as occlusion to depth order, and the lack of perspective information due to the relatively large viewing distances. The projections thus approximate parallel projections. The depth order/rotation direction of parallel projections can theoretically be disambiguated by occlusion, stereo and perspective cues.

Occlusion cues

Braunstein, Andersen and Riefer (1982) showed that occlusion could be used to resolve the ambiguity of depth order in a parallel projection of an object rotating in depth. The stimuli were parallel projections of pentagonal elements on the surface of a rotating sphere. In one condition, the sphere was treated as opaque: elements disappeared as they rounded the edge of the sphere (edge occlusion). In the other condition, the sphere was treated as transparent, with the elements on the near edge occluding those on the further edge (element occlusion). Subjects judged the direction of rotation. It was found that the former type of occlusion produced more accurate percepts than the latter. Accuracy increased with increasing element size. A follow-up study by Andersen and Braunstein (1983) demonstrated that dots located within implicit pentagonal texture elements were just as effective as disambiguating rotation direction, showing that visible contours were not necessary.

The relative abilities of the two types of occlusion, edge and element, to dominate stereo information in the perceived depth order was investigated by Braunstein, Andersen, Rouse and Tittle (1986). It was found that perceived rotation direction was dominated by edge occlusion in over 60% of trials, and was dominated by element occlusion in approximately 15% of trials. Translating dihedral angles, in which the velocity gradient specified a given depth order and stereo information specified the opposite, were perceived in accordance with the stereo information when this specified a 'centre near' surface (90%) more often than when it specified a 'centre far' surface (50%).

Occlusion information can also disambiguate the depth order of frontal surfaces which translate to and fro across the line of sight, both when the observer remains stationary (Kaplan, 1969) and when the parallax is linked to their head motion (Ono, Rogers, Ohmi and Ono, 1988). In the latter study, when the motion parallax and occlusion cues specified opposite depth orders, the perceived order was dependent upon the size of the parallax cue. When the simulated depth separation of the surfaces was small, the perceived depth order was determined by the parallax information. When the depth separation was large, the order was determined by the dynamic occlusion information. These results were

interpreted as suggesting that parallax information is most appropriate for specifying the 3-D distances *within* objects since these distances are generally small, whereas dynamic occlusion information is more suited to specifying the depth order *between* objects situated at different distances.

Stereo cues

In cases where motion information is ambiguous with respect to depth order and rotation direction, the addition of stereo cues can influence these percepts. Graham and Rogers (1982b; see Section 1.4.5) showed that the perceived depth order of an ambiguous motion parallax surface (essentially a KDE) could be influenced by previous exposure to a stereoscopic surface with unambiguous depth order. The perceived phase of the parallax surface was opposite to that of the adapting stereo surface.

A rotating sphere under parallel projection generally undergoes continuous spontaneous reversals in depth and concomitant rotation direction. Nawrot and Blake (1991) asked observers firstly to watch such a sphere for 10 minutes; during this time their perceptual reversals were recorded. After a 90-second period of adaptation to a sphere in which the rotation direction was disambiguated by stereoscopic information, observers viewed the ambiguous sphere once more. The perceived direction of rotation was opposite to that experienced during the stereoscopic adaptation period. These results complement previous studies in which observers who could not detect relative depths in a static disparity stimulus could determine the rotation direction of a sphere under parallel projection when disparity cues were added (Braunstein, Andersen, Rouse and Tittle, 1986; Rouse, Tittle and Braunstein, 1989). They also complement the results of an earlier study by Nawrot and Blake (1989), in which observers first adapted to a rotating sphere in which the dot disparities provided unambiguous information about the rotation direction and depth order, and then regarded an ambiguous rotating sphere in which the dots had zero disparity. The latter sphere was perceived to rotate in the opposite direction.

This evidence is suggestive of a strong interaction between stereopsis and motion in the perception of depth order and rotation direction in motion-defined structures. It complements other work demonstrating

interactions between these modalities (e.g. Graham and Rogers, 1982a; Rogers and Graham, 1984; Bradshaw and Rogers, 1993, 1996; see Section 1.4).

Vertical perspective cues

The strong effect of monocular perspective upon perceived depth order and rotation direction is provided by the rotating trapezoid illusion (Ames, 1951). When a perspective drawing of a window receding in depth was rotated about its vertical axis while viewed monocularly from a distance, it appeared to oscillate. When the longer side of the trapezium was nearer the observer, the correct direction of rotation was reported; when the shorter side was nearer, the apparent direction was reversed. A rectangular window rotated under the same conditions gave rise to consistently veridical perceptions of its rotary direction. Zegers (1964) found that the instances of reversal of the rotating trapezoid were decreased in frequency when the stimulus was large, and suggested that this might be due to the increased strength of the perspective information specifying the veridical direction of rotation.

As noted in Section 6.1, Gibson and Gibson (1957) demonstrated that vertical perspective was an effective cue to the slant angle of a surface, especially when the surface rotated to the given angle rather than being presented, stationary, at the specified slant. Surfaces with either a regular or irregular outline were used. Neither produced reversals in rotation direction, leading the authors to suggest that perspective information was not specifically carried in the outline shape of the stimuli. Power (1967) attempted to determine the relative effectiveness of two types of visual transformation in the determination of perceived rotation direction. He noted that both (i) the vertical height ratio of the vertical lines of a stimulus and (ii) the orientation of the horizontal lines of the stimulus both change systematically as the object rotates. Four types of stimuli were used. Type one contained a horizontal pair of lines whose ends were joined by irregular lines, type two contained a vertical pair of lines whose end were joined by irregular lines, type three contained only irregular edges and type four contained straight edges only. The stimuli rotated about either a vertical or a horizontal axis, and observers reported the occurrence of spontaneous reversals in perceived rotation direction.

Reversals were reduced when the stimulus contained straight edges which were perpendicular to the axis of rotation, thus Power (1967) proposed that the change in orientation of these lines was the more important means by which perspective exerts its influence. A subsequent experiment by Braunstein and Payne (1968) found supporting evidence: vertical perspective information was found to be more important in the determination of the rotation direction of rotating trapezoids. Börjesson (1971) also provides evidence to support Power (1967). By controlling the orientations and lengths of the edges of a trapezoid, he showed that the frequency of rotation direction reversals was reduced in the presence of appropriately converging horizontal edges. He also demonstrated the importance of textural perspective in the perception of the correct rotation direction. Braunstein (1972) developed a model for the perception of rotation in depth. The change in angle between the horizontal and vertical contours was held to be the main indicator of rotation direction, with the difference in acceleration of the vertical sides playing a secondary role. Predictions of the model were found to quantitatively match the results of Zegers (1964), Braunstein and Payne (1968) and Braunstein (1971).

Further evidence for the relative importance of vertical perspective is provided by Hershberger, Stewart and Laughlin (1976). These authors produced stimuli in which the horizontal perspective information conflicted with the vertical perspective information, and found that perceived rotation direction was dominated by the latter.

Braunstein (1962, 1977a) found that observers were increasingly likely to perceive three-dimensionality in motion displays when increases in polar perspective occurred. Braunstein (1977b) also attempted to determine the relative contributions of horizontal and vertical perspective to the accurate perception of depth order. Observers were presented with dot patterns simulating points randomly distributed in a sphere which rotated about a vertical axis. The displays were produced using either normal polar projections, or with perspective effects limited either to the horizontal or the vertical dimension. Two simulated distances and two display sizes were used. The subject judged the direction of rotation of the sphere. It was found that the proportion of accurate judgements in the normal polar projection and the vertical perspective only conditions were similar (around 0.90 and 0.75 for the nearer and further simulated

distances respectively), but was at chance in the horizontal perspective only condition. This study demonstrates the greater utility of vertical perspective as a cue to depth order — although note that the horizontal perspective effects in a rotating sphere are very small and might thus not be expected to strongly influence the perceived depth order.

Jansson and Börjesson (1969) presented subjects with a single vertical line whose horizontal position and vertical size were systematically varied so as to simulate a perspective projected rotation. When the same stimulus contained only velocity cues, observers could not make accurate rotation judgements. This suggests that motion information combined with a length change is a particularly effective carrier of perspective information.

Power and Day (1973) suggested that the perceived rotation direction of a structure will be unstable if its depth structure is incorrectly perceived. Subjects were presented with dihedral angles whose surface markings provided information either agreeing or disagreeing with the true convexity or concavity of the surface. Subjects viewed the rotating objects monocularly, and it was found that the perceived shape of the surface was determined by the surface markings. Rotation direction was determined by the perceived shape.

In a similar vein to the disambiguation of the depth order of parallel projections by unambiguous stereo information (Nawrot and Blake, 1991), disambiguation of a rotating sphere can also be provided by prior adaptation to a perspective simulation of a rotating square (Petersik, Shepard and Malsch, 1984). In sum, these studies provide evidence strongly suggestive of the importance of vertical perspective information in the determination of perceived rotation direction and depth order.

Despite the fact that perspective projections are theoretically unambiguous with respect to depth order, reversals have been reported. The ambiguities in depth order experienced by stationary observers viewing displays of translating surfaces (Braunstein and Andersen, 1981; Farber and McConkie, 1979; Braunstein and Tittle, 1993) contrast greatly with the lack of ambiguity in displays in which motion parallax is linked either to the motion of the observer (e.g. Rogers and Graham, 1979, 1982;

Ono et al., 1986; Ono and Steinbach, 1990) or to the display scope (Rogers and Graham, 1979; Rogers and Collett, 1989). A common feature of the latter two paradigms is the relative motion between the entire flow field and the observer's head, and the necessity of this feature for unambiguous depth perception was investigated by Braunstein and Tittle (1988). In one condition, the parallax surface was placed in a window, which translated to-and-fro across the screen. In another condition, the window remained stationary but the observer-relative motions in the parallax surface were the same as those of the previous condition. It was found that the accuracy of perceived depth order was identical in these conditions, and the authors concluded that effective motion parallax does not require relative motion between the observer's head and the contours of the flow field. The fact that this is the case should not, however, be taken to indicate that other sources do not also contribute to the lack of ambiguity of parallax surfaces viewed in situations involving relative motion between the observer and the entire display. Rogers and Rogers (1992) decomposed the information arising in these situations into visual and non-visual factors. Examples of the latter include proprioceptive and vestibular signals produced during the observer's movement. The former factor refers to the optic flow information in the whole optic array. Foreground flow, real trapezoidal transformations and the proprioceptive and vestibular information were found to be particularly effective at disambiguating depth order. Vertical perspective changes were found to be as effective when presented alone as when accompanied by a horizontal width change of the stimulus, indicating the relative ineffectiveness of the latter.

In sum these studies illustrate the importance of vertical perspective information in specifying the perceived depth order and rotation direction, in stimuli viewed by both stationary and moving observers. Other studies have demonstrated its importance in determining the perceived magnitude of rotation (see Section 6.1).

Although these studies demonstrate the importance of vertical perspective in the determination of the perceived direction of rotation, they do not provide information concerning the extent to which its influence is dependent upon its magnitude.

In situations in which the relative motion creates the impression of a

corrugated surface — such as the ambiguous parallax paradigm used by Rogers and Rogers (1992) — the depth order is indeterminate. The central corrugation can be legitimately perceived as either a peak or a trough. As noted above (Section 6.8.1), the stimulus is essentially a KDE. With the addition of vertical perspective information only one of the two solutions (either 'central peak' or 'central trough') is consistent with the perspective cue.

The present experiment investigates the effect of rotation magnitude upon the perceived depth order. Assuming that the simulated rotation angle is supra-threshold, it might be the case that the proportion of time spent perceiving a given depth order is dependent only upon the *direction* of rotation signalled by the vertical perspective information. Alternatively, it may be the case that the *magnitude* of simulated rotation also plays a role.

6.8.2 Stimuli

Stimuli were identical to those described in previous experiments, except that the relative dot motion was externally generated and the viewer remained stationary. The frequency and amplitude of the relative motion were 0.5 Hz and 20 arc min equivalent disparity respectively. Stimulus size was 5° or 20°. The vertical perspective cue had five levels. These simulated rotation angles of 0°, ±2° and ±10°. Positively-signed angles indicate that the vertical information was consistent with a central peak. Negatively-signed angles indicate that the perspective information was consistent with a central trough. In the 0° condition the depth order was ambiguous.

6.8.3 Procedure

Each of the 10 conditions (two sizes x five vertical perspective levels) was carried out 6 times. The order of the conditions was randomised.

Trials lasted 30 seconds. During this time subjects pressed one of two buttons according to whether the central horizontal corrugation was perceived as a peak or a trough. If it was perceived as neither, the subject pressed neither button. The perceived depth order was monitored

continuously throughout each trial. Subjects were free to move their eyes around the stimulus, but were instructed to make a judgement about the appearance of the central corrugation.

6.8.4 Results

Figs. 6.61 - 5.63 show the Mean Proportion of Response Time per trial in which the central corrugation was perceived as a Peak (black bars), a Trough (white bars) or as Non 3-D (grey bars) for each rotation angle simulated by the vertical perspective cue, for Subjects SAS, SJP and MML respectively. The bars to the left of the vertical dashed line represent the data for the size 5° stimulus. The bars to the right of the vertical dashed line represent the data for the size 20° stimulus. Fig. 6.64 shows the data averaged over subjects.

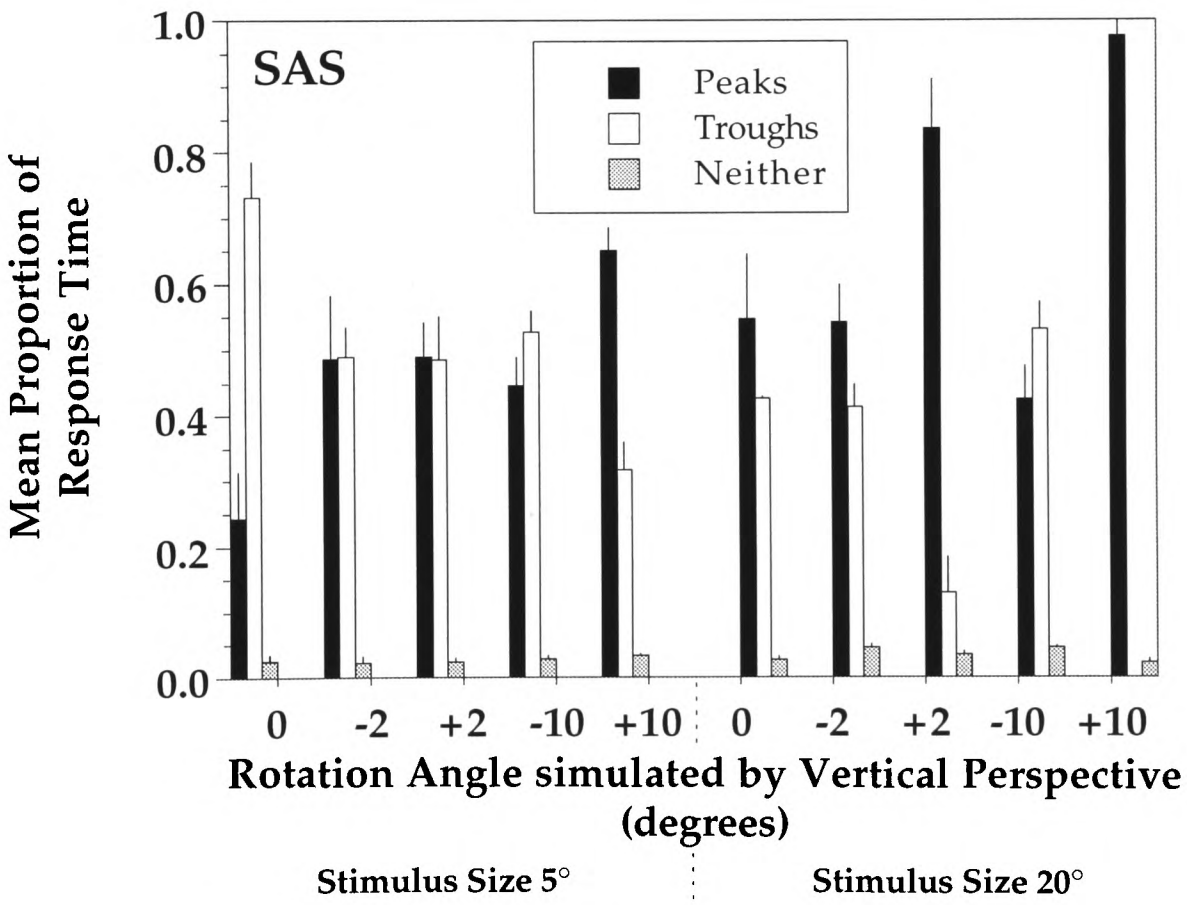


Fig. 6.61: Mean Proportion of Response Time during which the central corrugation was perceived as a Peak (black columns), a Trough (white columns) or Neither (grey columns) for each rotation angle simulated by the vertical perspective cue, for Subject SAS. The bars to the left of the vertical line represent data for the size 5° stimulus, and the bars to the right of the line represent data for the size 20° stimulus.

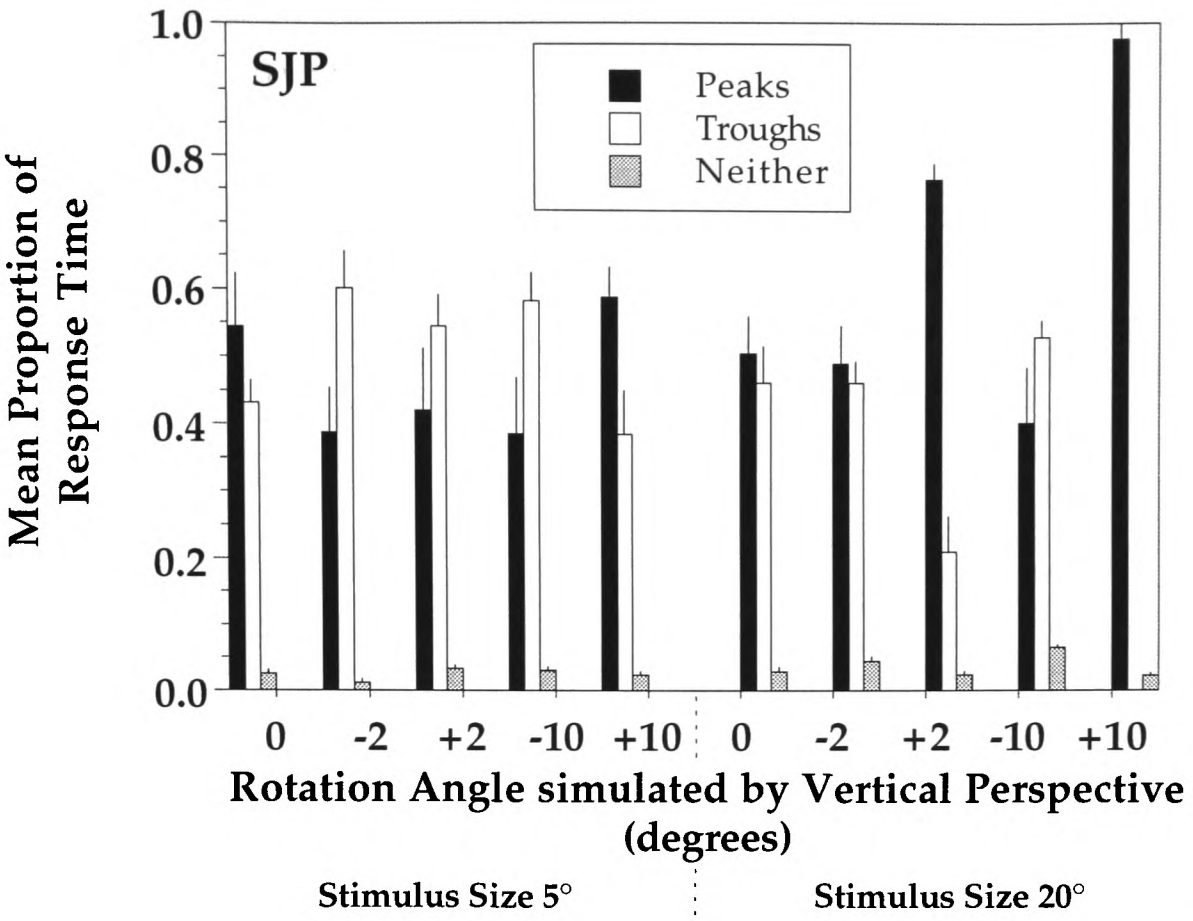


Fig. 6.62: Mean Proportion of Response Time during which the central corrugation was perceived as a Peak (black columns), a Trough (white columns) or Neither (grey columns) for each rotation angle simulated by the vertical perspective cue, for Subject SJP. The bars to the left of the vertical line represent data for the size 5° stimulus, and the bars to the right of the line represent data for the size 20° stimulus.

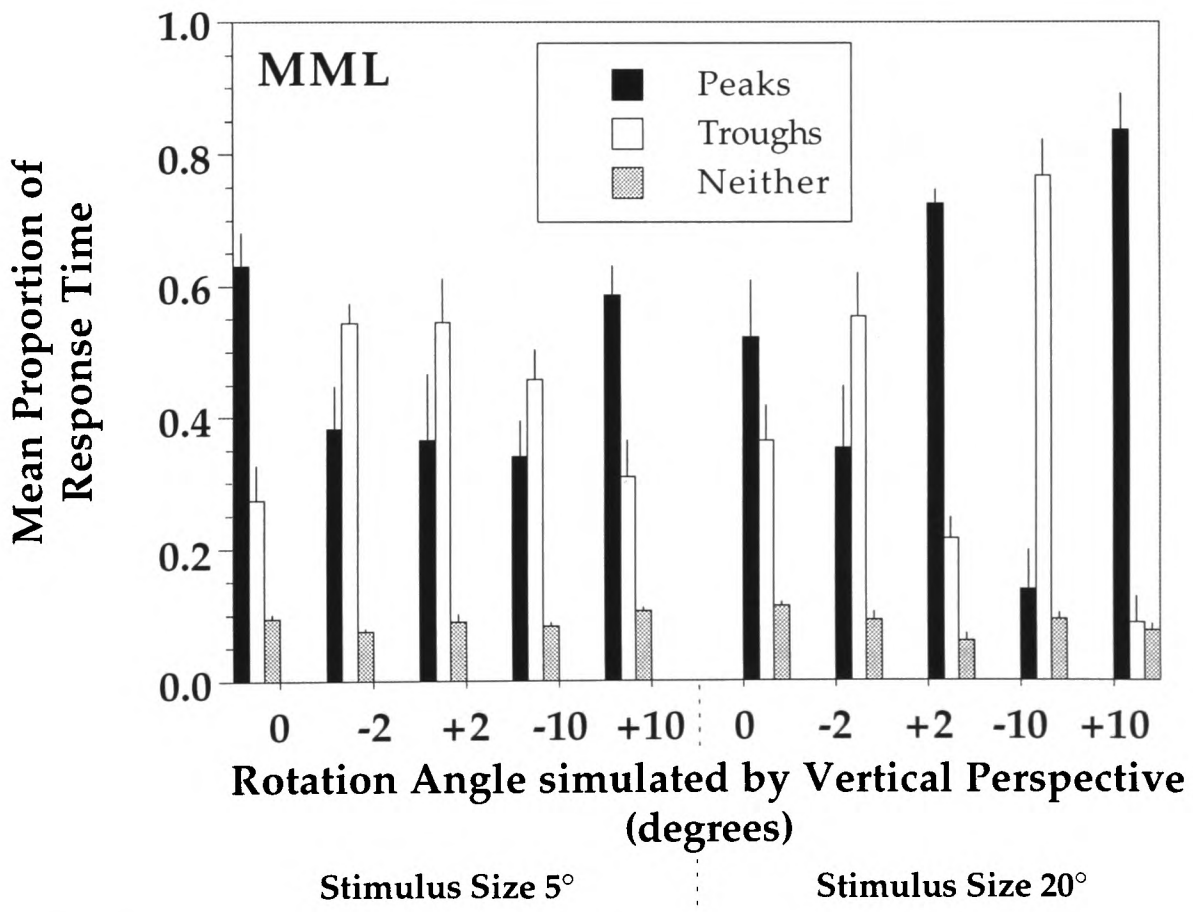


Fig. 6.63: Mean Proportion of Response Time during which the central corrugation was perceived as a Peak (black columns), a Trough (white columns) or Neither (grey columns) for each rotation angle simulated by the vertical perspective cue, for Subject MML. The bars to the left of the vertical line represent data for the size 5° stimulus, and the bars to the right of the line represent data for the size 20° stimulus.

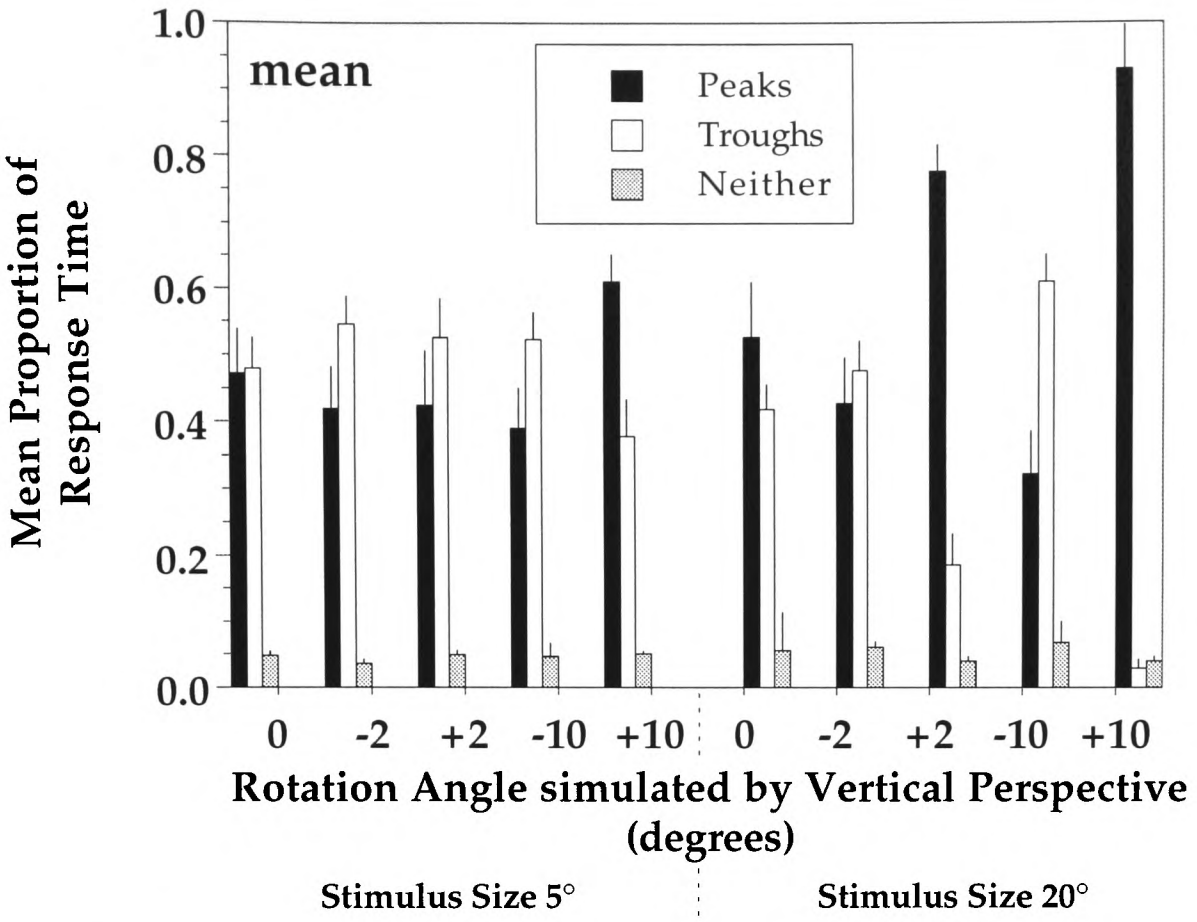


Fig. 6.64: Mean Proportion of Response Time during which the central corrugation was perceived as a Peak (black columns), a Trough (white columns) or Neither (grey columns) for each rotation angle simulated by the vertical perspective cue, averaged over three subjects. The bars to the left of the vertical line represent data for the size 5° stimulus, and the bars to the right of the line represent data for the size 20° stimulus.

When vertical perspective simulated a rotation angle that was consistent with a central peak, the proportion of peaks perceived was generally greater than the proportion of troughs. When the rotation angle was consistent with a central trough, the proportion of troughs perceived was greater than the proportion of peaks. This effect was more marked in the 20° stimulus, and increased with the magnitude of simulated rotation such that in the ± 10° conditions the depth order was almost completely dominated by the vertical perspective cue.

A two-way repeated-measures ANOVA was carried out upon these pooled data, with size (2, 20°) and simulated rotation angle (0, ±2, ±10°) as the factors. Subject was the repeated factor. The dependent variable was the mean proportion of time spent perceiving a peak, henceforth referred to as Peak Time. Inspection of Figs. 6.60 - 6.64 reveals that the mean proportion of response time spent perceiving neither a peak nor a trough is similar across conditions. Peak Time and Trough Time are thus inter-dependent; either variable could equally well be entered into the analysis as the dependent variable.

The main effect of size was non-significant ($F[1, 2] = 10.14$, n.s.). The main effect of simulated rotation angle was significant ($F[4, 8] = 13.13$, $p<0.01$). The interaction between size and simulated rotation angle was significant ($F[4, 8] = 9.92$, $p<0.005$).

Tukey's HSD tests were employed to determine the significance of the differences between the Peak Times of the simulated rotation angle levels at each size level.

$$HSD = q_{(0.05)}(p, v) * \sqrt{(MS_{error}/n)}$$

Where
n = number of observations in each mean
p = number of means
v = p(n-1).

n = 3
p = 5
v = 5(3-1) = 10.

$q_{(0.05)}(5, 10) = 4.65$

$$\therefore HSD = 4.65 * \sqrt{(0.01/3)}$$
$$= 0.27$$

Size 5°

simulated rotation angle		**-10°**	**-2°**	**+2°**	**0°**	**+10°**
	Peak Time	0.40	0.42	0.43	0.47	0.53
-10°	0.40	*0.00*				
-2°	0.42	*0.02*	*0.00*			
+2°	0.43	*0.03*	*0.01*	*0.00*		
0°	0.47	*0.07*	*0.05*	*0.04*	*0.00*	
+10°	0.53	*0.13*	*0.11*	*0.10*	*0.06*	*0.00*

Table 6.1: Peak Time (second row and second column) for each simulated rotation angle, for stimuli of size 5°. Each italicised number represents a difference score. Significant (at the 0.05 level) differences are marked with an asterisk.

Size 20°

simulated rotation angle		**-10°**	**-2°**	**0°**	**+2°**	**+10°**
	Peak Time	0.32	0.46	0.53	0.78	0.93
-10°	0.32	*0.00*				
-2°	0.46	*0.14*	*0.00*			
0°	0.53	*0.21*	*0.07*	*0.00*		
+2°	0.78	*0.46**	*0.32**	*0.25*	*0.00*	
+10°	0.93	*0.61**	*0.47**	*0.60**	*0.15*	*0.00*

Table 6.2: Peak Time (second row and second column) for each simulated rotation angle, for stimuli of size 20°. Each italicised number represents a difference score. Significant (at the 0.05 level) differences are marked with an asterisk.

Despite the non-significant main effect of size, the interaction of size and simulated rotation angle was significant. Thus Tukey's HSD tests were again employed to determine the significance of the differences between Peak Times of the size levels at each simulated rotation angle.

$n = 3$

$p = 2$

$v = 2(3-1) = 4.$

$q_{(0.05)}\ (2, 4) = 3.93$

$\therefore$ HSD $= 3.93 * \sqrt{(0.01/3)}$

$= 0.23.$

Table 6.3 shows the Peak Times for each combination of simulated rotation angle and stimulus size. Each italicised number represents the difference between the two preceding numbers of the column (i.e. represents the difference due to stimulus size). Significant differences are marked with an asterisk.

	simulated rotation angle	**0°**	**-2°**	**+2°**	**-10°**	**+10°**
stimulus size						
5°		0.47	0.42	0.43	0.40	0.53
20°		0.53	0.46	0.78	0.32	0.93
difference		*0.06*	*0.04*	*0.35**	*0.12*	*0.40**

Table 6.3: Peak Times are shown in plain text. The italicised numbers represent the difference between the two preceding numbers of the column. Significant (at the 0.05 level) differences are marked with an asterisk.

These results show that the influence of the simulated rotation angle upon the amount of time spent perceiving a given depth order is crucially dependent upon the stimulus size. Examination of Table 6.1 reveals that there is a non-significant effect of simulated rotation angle for stimuli of size 5°. Conversely, a number of Peak Times of stimuli of size 20° are significantly different from each other.

Table 6.2 shows that the largest difference between Peak Times is between those of the +10° and -10° surfaces. Table 6.2 also shows that the difference between the Peak Times of the surfaces of 0° and +10° simulated rotation is greater (and is significant) than the difference between the respective Peak Times of the surfaces of simulated rotation angle 0° and -10° (which is non-significant). Likewise, the difference between the Peak Time of the surface of 0° simulated rotation and that of the surface of +2° simulated rotation is greater than the difference between the respective Peak Times of the surfaces of simulated rotation angle 0° and -2°. This suggests that, given that the stimulus size is large, a dynamic perspective cue may be more influential in determining perceived depth order when the cue specifies a peak than when it specifies a trough. Furthermore, Table 6.3 shows that the effect of stimulus size is significant only when the simulated angle of rotation is +2° or +10°. This is due to the much elevated Peak Times of the larger stimuli of these rotation angles. The effect provides further evidence that the dynamic perspective cue is particularly effective at influencing perceived depth order when the specified order is positive.

6.8.5 Discussion

As outlined earlier, there is ample evidence to suggest that the perceived rotation direction and accompanying depth order of a stimulus are strongly influenced by a dynamic vertical perspective cue (e.g. Ames, 1951; Zegers, 1964; Gibson and Gibson, 1957; Power, 1967; Braunstein and Payne, 1968; Börjesson, 1971; Power and Day, 1973; Hershberger et al., 1976; Braunstein, 1977; Rogers and Rogers, 1992). The current experiment illustrates that the influence of the perspective information generally increases with increasing simulated rotation angle and with increasing stimulus size. Experiment 17 of this chapter demonstrated that for a given (non-zero) magnitude of simulated rotation, the magnitude of perceived rotation increases with increasing stimulus size. The present experiment has shown that the direction of rotation can influence the depth order increasingly strongly when stimulus size is large. However, it should not automatically be concluded that this effect occurs *because* the magnitude of perceived rotation is greater. Petersik (1980) found that an increase in the magnitude of perceived rotation of a sphere went hand in hand with a decrease in the accuracy with which its depth order was perceived. It may be the case that the effect would still occur when the magnitudes of perceived rotation of small and large stimuli are equal. A way of testing this would be to reduce the amount of vertical perspective in the large stimulus until the two sizes appeared to rotate to the same extent, and then to examine perceived depth order.

Inspection of the graphs of the individual subjects reveals that a slightly greater proportion of peaks than troughs was perceived in conditions in which depth order was ambiguous. This result is similar to that found by Hayashibe (1991), and is dissimilar to the result of Rogers and Rogers (1992), who found that in their ambiguous condition subjects experienced many reversals in depth, and that ".. within any trial the target corrugation of the surface appeared to be concave as often as it appeared convex".

The present result may be partially explicable by the results of Hayashibe's Experiment 3. Subjects were presented with square wave surfaces and were asked to make a depth order judgement. They were then instructed to fixate the strips which appeared to recede and to report

whether these strips then underwent depth reversal. It was found that subjects generally fixated the strips which had the slower velocity, and that these initially appeared to protrude, but that when fixation moved to the apparently receding strips these underwent depth reversal. When fixation was maintained on protruding strips, no reversals occurred. In the present experiment it may be the case that, despite the fact that subjects were free to move their eyes around the stimulus, since they were instructed to make their judgements on the depth order of the central corrugation they generally tended to fixate this corrugation. If it was perceived as a peak it would be expected, on the basis of Hayashibe's findings, that it would remain a peak in the absence of shifts in fixation to an apparent trough. Reversals may have occurred when fixation was shifted to an adjacent trough.

6.9 Summary and Conclusions

The experiments of Chapter 5 demonstrated that a trade-off relationship exists between the perceived extents of depth and rotation when either corrugation frequency or stimulus size is manipulated. The present chapter aimed to determine whether the trade-off relationship exists when a dynamic vertical perspective cue was used to manipulate the perceived rotation magnitude. Experiments 16 and 17 showed that the relationship existed only for size 20° surfaces of relatively high corrugation frequency. The size of the trade-off was much smaller than that observed in the previous chapter. Although the comparisons which may be drawn between the results of the two chapters are limited because of the different depth matching techniques employed, the data suggest that it is not necessarily manipulations of rotation *per se* which lead to alterations in the perceived extent of depth. Rather, the extent of the trade-off relationship is dependent upon the nature of the manipulated stimulus variables.

In Experiment 18, subjects adjusted the relative motion amplitude of a triangle wave surface so that the dihedral angle between the two sides of the peaks (and troughs) appeared to be 90°. Experiment 19 involved the same judgement, but observers adjusted the corrugation frequency of the surface. This task was performed for various simulated rotation angles.

When the simulated rotation direction was in the *opposite* direction to the head movement of the observer, the relative motion amplitude required for the surface to contain perceived right-angles increased with increasing simulated rotation angle. These results are qualitatively consistent with a depth/rotation trade-off relationship. When the simulated rotation direction was in the *same* direction as the head movement, the surface was perceived to rotate in the opposite direction, i.e. the percept was inconsistent with the vertical perspective information. No clear pattern of perceived depth magnitudes was evident. The propensity of the observers to perceive rotation in the opposite direction to the head movement in the face of simulated rotation in the same direction as the head movement is consistent with a bias of the visual system to interpret elements moving in the opposite direction to the head motion as counter-rotating, even at the expense of non-rigidity. It was suggested that a further factor contributing to this putative bias is the real-world vertical perspective information arising from the display upon observer movement. This information specifies a stationary, frontal surface. With respect to the eye, however, the surface rotates in the opposite direction to the head movement. A negative depth order is geometrically inconsistent with this rotation direction (see Fig. 5.6).

Experiment 20 examined the effect of vertical perspective on perceived rotation magnitude as a function of viewing distance. Surfaces with identical simulated rotation angles were found to rotate more as viewing distance increased, but the difference between the magnitudes of perceived rotation at two vertical perspective levels at a given distance generally remained constant. This was interpreted to mean that, under these conditions at least, the weighting of the vertical perspective information in determining the perceived rotation magnitude remains approximately constant over viewing distance. It was suggested that the results are instead consistent with a propensity of the visual system to increasingly attribute, with increasing viewing distance, vertical perspective changes to object rotation rather than observer translation.

In Experiment 21, observers remained stationary and the motion parallax was externally generated. The experiment examined the ability of dynamic vertical perspective to influence the perceived depth order in an otherwise ambiguous stimulus. It was demonstrated that the proportion

of time spent perceiving a given depth order is dependent both upon the magnitude of simulated rotation and upon the size of the stimulus.

Chapter 7:

Conclusions

7.1 Thesis Aims

This thesis aimed to provide insights, both qualitative and quantitative, into the mechanisms underlying the perception of surfaces defined by motion parallax. The thesis was motivated by the fact that motion parallax is often considered to be an ambiguous cue, since a given pattern of retinal image motions can theoretically arise from a number of different scenarios. In the real world there are multiple sources of information about the structure and movement of objects in the visual environment. The information provided by motion parallax is accompanied by various other cues whose incorporation into the percept severely constrains the number of alternatives open to the system. Richards (1985), for instance, proposed a model in which the shape of an object is extracted from motion information, while binocular disparity specifies the sign of depth. Numerous psychophysical studies have demonstrated that the interpretation of parallax information is influenced by binocular disparity (e.g. Braunstein et al., 1986; Rouse et al., 1989; Rogers and Collett, 1989; Nawrot and Blake, 1991), occlusion cues (e.g. Braunstein et al., 1982, 1986; Ono et al., 1988) and perspective information (Gibson and Gibson, 1957; Power, 1967; Braunstein and Payne, 1968; Power and Day, 1973; Braunstein, 1977).

The current work did not employ such additional cues, since the aim was to isolate the information provided by motion parallax alone. Faced solely with such information the visual system must select, from the rich variety of possible interpretations of a given stimulus, a single interpretation. The aim was to examine the various factors which govern the interpretation as revealed by the observer's perceptions.

The majority of the experiments involved a head-linked motion

parallax paradigm, in which the relative motion of the display dots was linked to the movement of the observer. Much of the existing literature assumes that the relative dot motion in this situation is interpreted as arising from a stationary, 3-D structure. However, the fact that a variety of different scenes can theoretically give rise to the same optical transformation means that legitimate solutions involving the global motion (a rotation, for example) of a slightly different 3-D structure also exist.

It has been noted by a number of studies that a rotary component is frequently perceived in structures defined by motion parallax (e.g. Braunstein and Andersen, 1981; Ono et al., 1986; Rogers and Collett, 1989; Ono and Steinbach, 1990). Other studies have noted that, under different circumstances, a 'shearing' motion is perceived (Ono and Steinbach, 1990; Ono and Ujike, 1993; Ichikawa and Ono, 1996). However, there is a paucity of work providing a systematic, quantitative analysis of the qualitatively different percepts to which a parallax stimulus may give rise. This thesis attempted to go some way towards providing such an analysis.

7.2 Summary

Experiments 1 to 20 used head-linked motion parallax surfaces, and investigated observers' percepts of depth, rotation and non-rigidity. Experiment 21 investigated the perceived depth order of surfaces in which the parallax was externally generated and the observer remained stationary.

7.2.1 Chapter 3

A small number of experimenters has suggested that different parallax surfaces can fall into different 'perceptual zones' which are divided by perceptual thresholds. For instance, Ono and Ujike (1993) and Ichikawa and Ono (1996) noted that observers perceived a square wave motion parallax surface to undergo shearing motion under some conditions and to remain stationary under others. Investigating the point at which the shearing started to occur (the 'motion threshold') as a function of head velocity, they found that as head velocity increased, the threshold

decreased with a slope of approximately -1. It was suggested that the motion parallax was 'fully effective' (i.e. was interpreted solely as depth) in the region between the depth threshold and the 'motion threshold'. At head velocities greater than the 'motion threshold', motion parallax was not fully effective. The nature of the motion perceived above the 'motion threshold' has been shown to be dependent upon the nature of the waveform (Ono and Steinbach, 1990). Sinusoidally depth modulated surfaces appear to rotate under certain circumstances (e.g. Ono et al., 1986; Ono and Steinbach, 1990). The phenomenon has also been observed under conditions in which a dihedral angle (a horizontally oriented 'wedge') translates with respect to a stationary observer (Braunstein and Andersen, 1981). In this study, a rotary component was perceived in the display, despite the fact that none was simulated.

Quantitative measurements of the point at which sinusoidal motion parallax surfaces start to appear to rotate have not previously been made. In Chapter 3, the relative motion amplitude at which rotation is first perceived was termed the 'Transition Point'. The Transition Point represents the border between two different perceptual interpretations of the parallax stimulus. One interpretation is of a surface which is rigid and stationary; the other interpretation is of a surface which is rigid and rotating. It was suggested that the border between these percepts might be rather conceptual in nature, since there exists no dichotomy of perceptions but rather a smooth continuum. Nonetheless, subsequent experiments explored the way in which a number of variables affected the Transition Point.

The experiments of Chapter 3 employed the Method of Constant Stimuli to investigate the Transition Point as a function of a variety of factors.

Experiment 1a determined the Transition Point as a function of corrugation frequency. It was shown that over a four octave range (0.05 to 0.8 cpd) the Transition Point was primarily determined by the maximum motion gradient (MMG) of the surface such that, when plotted against corrugation frequency, the functions had slopes which were close to -1. At frequencies greater than 0.8 cpd (0.4 cpd for one subject) the function levelled off, suggesting that at this point Transition Points became

dependent upon the absolute relative motion amplitude. Using 'clipped' surfaces in Experiment 1b — created by doubling or quadrupling the amplitude of the unclipped waveform and clipping the peaks and troughs of the resulting waveform at 50% or 25% of their respective heights — the role of the MMG was further investigated. Surfaces of identical MMG were generally found to have similar Transition Points, providing evidence in accord with the notion that the motion gradient is the prime determinant of the Transition Point over a range of frequencies. Also consonant with the notion was the observation that, for a given corrugation frequency, Transition Points fell with increases in the degree of 'clippedness'. The Transition Points of clipped stimuli were sometimes greater than predicted, and it was suggested that this might be due to the perception of a slight non-rigidity of the surface at the expense of its rotation, such has been reported with square wave corrugations (Ono and Steinbach, 1990; Ono and Ujike, 1993; Ichikawa and Ono, 1996).

Experiment 2 investigated Transition Points for four different relative motion gradients: sine, triangle, ramp and square waveforms. Square waveforms were shown to have the highest Transition Points, followed by ramp, sine and triangle waveforms, in that order. The high Transition Point of the square wave surface indicated that rotation was perceived least readily for this waveform. This was attributed to observers' propensity to perceive 2-D shearing motion in preference to rotation. No evidence was found to suggest that the visual system analysed the surfaces on the basis of their respective Fourier components. Rather, the similarity of the Transition Point functions of the sine and triangle waveforms suggested that the motion gradients played a more important role than the Fourier components in determining the Transition Point of the latter surface at least. This result accords with those of Experiments 1a, 1b and 2, and with studies by Welch (1996) and Sachtler and Zaidi (1995), which also point to the importance of the motion gradient in observers' perception of compound motion stimuli. Additionally, it has been demonstrated that the perception of stereoscopic surfaces is determined primarily by surface curvature rather than Fourier components (Norman et al., 1991). It was suggested that the Transition Point functions of the square and ramp waveforms were best accounted for by an analysis of both the maximal and non-maximal motion gradients.

Experiment 3 examined the effect of stimulus size on the Transition Point, and demonstrated that Transition Points decrease with decreasing stimulus size. It was suggested that this was due to the decreasing strength of the conflicting vertical perspective information specifying a lack of rotation with decreasing stimulus size. The gradient of the Transition Point function was substantially shallower than the gradients of the functions determined in Experiment 1a, perhaps indicating a decreased importance of the vertical perspective cue (relative to other cues) when the rotation is at threshold levels.

Experiment 4 was motivated by the observation of Ono et al. (1986) that the number of instances of perceived rotation of a parallax surface increased with increasing viewing distance. It would therefore be predicted that the Transition Point should fall with increasing distance. Ono et al. (1986) used the same distal stimulus throughout, such that increases in the viewing distance also led to an increase in surface corrugation frequency and a decrease in the proximal stimulus size. Experiments 1a and 3 showed that both of these factors lead to decreases in the Transition Point. Using stimuli which were proximally identical in order to eliminate these potential confounds, Experiment 4 showed that the Transition Point decreased as the viewing distance increased from 57 to 228 cm — i.e. subjects are increasingly likely to perceive rotation as the viewing distance increases. This demonstrates that the threshold is not quantifiable merely in terms of the proximal pattern of image motion. It was suggested that the importance of the vertical perspective information in specifying a stationary, frontal surface might decrease with increasing viewing distance, thus facilitating the perception of rotation.

Experiment 5 investigated the zones into which the percepts arising from a motion parallax stimulus may fall. The Transition Point, as defined above, divides two such zones: one in which the percept is of a rigid, stationary surface and another in which a rigid, rotating surface is perceived. A further zone was identified by Saida and Ono (1989) — at high relative motion amplitudes (>60 arc min equivalent disparity), the perceived depth order of the parallax surface became unstable, the magnitude of perceived depth decreased and 2-D shearing motion was perceived. Ono et al. (1986) found that at large viewing distances, observers perceived only 2-D motion in a parallax display, suggesting the existence of

yet another perceptual zone: one in which the proximal stimulus is directly perceived rather than being interpreted as arising from a 3-D structure.

Despite the reports of the various percepts which arise from parallax stimuli, rigorous work of a quantitative nature has generally not been carried out. Moreover, no *single* experiment to date has mapped the various perceptual zones into which a parallax stimulus can fall. Experiment 5 determined four perceptual thresholds as a function of corrugation frequency, in order to map five perceptual zones. These zones were, first - a zone in which no depth and no motion was perceived (the zone beneath the depth threshold), second - a zone in which a rigid, stationary surface was perceived (the zone between the depth threshold and the Transition Point), third - a zone in which a rigid, rotating surface was perceived (the zone between the Transition Point and the Upper Transition Point), fourth - a zone in which a non-rigid surface was perceived (the zone between the Upper Transition Point and the 2-D Shearing Threshold and fifth - a zone in which only 2-D motion was perceived (the zone above the 2-D Shearing Threshold).

The range of these zones was found to be dependent upon the corrugation frequency. The range of relative motion amplitudes supporting the percept of a rigid, stationary surface was much decreased at the higher corrugation frequencies. When the image motion cannot be explained by the perceived depth and/or rotation of a structure, the percept is of non-rigidity. The Upper Transition Point — the point at which a rigid, rotation surface becomes a non-rigid, rotating surface — was found to fall with increasing corrugation frequency at approximately the same rate as the Transition Point. This suggests that the processes underlying these qualitatively different percepts might possibly have something in common. In a sense, since non-rigidity involves high magnitudes of local rotation of the surface, the percepts should perhaps be thought of as existing at different positions on the same continuum. 2-D shearing thresholds were found to exceed the maximum relative motion amplitude provided by the experimental apparatus, meaning that the most extensive perceptual zone is that which contains non-rigid surfaces.

7.2.2 Chapter 4

The experiments of Chapter 4 were designed to build upon those of the previous chapter by examining perceived rotation in the perceptual zone above the Transition Point.

Experiment 6 compared the perceived rotation of two surfaces of different corrugation frequency. It was found that the surface of higher frequency was always perceived to rotate to a greater extent than the lower frequency surface. However, the functions relating the Rotation Match to the corrugation frequency of the Test surface were shallower than predicted if the rotation magnitude were dependent solely upon the MMG.

The remainder of the experiments reported in Chapter 4 employed a new matching technique to measure perceived rotation. This was developed in order that a direct estimate of, rather than merely a comparison between, perceived rotation magnitudes could be made. The magnitude of perceived rotation of a real paddle was matched to the perceived rotation of the test stimulus. Experiment 7 examined the effect of corrugation frequency manipulations upon the magnitude of perceived rotation. It was demonstrated that, for a given relative motion amplitude, a greater extent of rotation was perceived with increasing corrugation frequency. This result is consistent with the finding that the Transition Point decreases with increasing corrugation frequency (Experiments 1a and 1b) and the finding that surfaces of high corrugation frequency are perceived to rotate to a larger extent than surfaces of lower corrugation frequency (Experiment 6). When replotted in terms of the MMG of each corrugation frequency at each relative motion amplitude, the slope of the matched rotation function increased with decreasing corrugation frequency. These results suggest that there is a decreased reliance of the visual system upon the MMG in the perception of rotation at supra-threshold levels. This contrasts with the importance of the gradient at threshold rotation levels: the Transition Point was shown to be principally determined by the MMG in the parallax stimulus. Experiment 7b sought to further examine the role of the MMG in the perception of rotation by using clipped surfaces identical to those of Experiment 1b. If the magnitude of perceived rotation were dependent upon the motion

gradients of the surface then it would be predicted that surfaces of similar gradient would appear to rotate to similar extents. In fact it was found that, within groups of surfaces containing identical MMGs, perceived rotation generally increased with increasing corrugation frequency. Moreover, for 0.1 or 0.2 cpd surfaces, increasing the MMG via the clipping procedure did *not* lead to increases in the amount of perceived rotation. Slight increases in the rotation magnitudes of 0.05 cpd surfaces were observed when the motion gradients were increased. It was suggested that, for the 0.1 and 0.2 cpd surface, the maximum ('ceiling') level of perceived rotation which had been reached in the current experiment might exist due to the perception of surface non-rigidity. On the basis of the Upper Transition Point function determined in Experiment 5, it might be expected that a greater amount of non-rigidity would be perceived in surfaces of higher corrugation frequency, and that this non-rigidity might be perceived at the expense of rotation. Measures of non-rigidity were not taken in the current experiment, but the data of Williams (1993) suggests that under certain circumstances non-rigidity is indeed perceived at the expense of rotation. The results of the present experiment demonstrated that, for the corrugation frequency range within which the Transition Point is primarily determined by the MMG, the magnitude of perceived rotation at supra-threshold levels is not.

Braunstein and Andersen (1981) found no difference between the extent of perceived rotation of a translating dihedral angle when the central portion pointed towards the observer and when it pointed away from the observer. However, the method of measurement was comparatively crude and thus may not have uncovered small differences in the perceived rotation of the two structures. Experiment 8 compared the magnitude of perceived rotation of surfaces containing a central peak with those containing a central trough using the more sensitive matching method noted above. Surfaces of 0.05 cpd were perceived to rotate to a greater extent when the central corrugation was a peak than when it was a trough. This was not found to be the case for the 0.1 and 0.2 cpd surfaces. The difference between the magnitudes of rotation perceived in the 'central peak' and 'central trough' conditions decreased with increasing corrugation frequency. The qualitative difference between the results of the 0.05 cpd surface and those of the 0.1 and 0.2 cpd surfaces is unclear, but it was suggested that the decrease in the difference between perceived

rotation in the 'central peak' and 'central trough' conditions might have partly been due to the increasing proximity (with increasing corrugation frequency) of the neighbouring corrugations to the attended-to central corrugation — and thus perhaps the increased influence of these flanking corrugations upon the magnitude of perceived rotation.

Experiment 9 investigated perceived rotation magnitude as a function of stimulus size. On the basis of Experiment 3, which showed that the Transition Point increases with increasing stimulus size, it was predicted that perceived rotation magnitude would decrease with increasing stimulus size. This was found to be the case. As in the former experiment, the basis of this effect is likely to be the increase, with increase in stimulus size, in the magnitude of the vertical perspective cues constraining the surface to remain in the frontal plane, thereby discouraging the perception of rotation. Interestingly, the effect of stimulus size was found to be stronger in the current experiment than in the corresponding Transition Point experiment. This would suggest that vertical perspective plays a more important role in the determination of supra-threshold rotation than in the determination of threshold rotation levels. These results contrast with the relative importance of the MMG in the two situations. The gradient was shown to play a significant role in the determination of the Transition Point, but is less important at supra-threshold rotation levels.

Experiment 10 explored the perceived rotation of surfaces which contained two different corrugation frequencies, and examined the relative importance of the central and flanking regions of the surface in the determination of rotation magnitude. It was found, as expected, that surfaces of higher mean corrugation frequency were perceived to rotate to a greater extent than surfaces of lower mean frequency, irrespective of whether the higher frequency corrugations were in the central band or the peripheral flanks. The difference in perceived rotation magnitude between stimuli containing the same frequency in the central band and different frequencies in the flanks was less than the difference in perceived rotation magnitude between stimuli containing the same frequency in the flanks and different frequencies in the central band, indicating the relative importance of the central region in determining the extent of perceived rotation.

7.2.3 Chapter 5

Chapter 5 reported the results of several previous studies which found evidence for a 'trade-off' relationship in depth-from-motion stimuli. It noted, however, that much of the work was not founded upon explicit geometrical considerations. This is especially true for the head-linked motion parallax literature. This chapter introduced two geometric models which generate quantitative hypotheses relating the theoretical magnitudes of the attributes involved in the 'trade-off' relationship. The first model provides an analysis of the covariation of depth and rotation magnitudes which theoretically exists for a given amount of relative motion within sinusoidally depth-modulated surfaces. The second provides an analysis of the covariation of depth and shearing magnitudes which theoretically exists for a given amount of relative motion within square wave depth-modulated surfaces. The former model was a modified version of that developed and tested by Williams (1993).

The experiments of the current chapter measured the perceived depth and rotation magnitudes of sinusoidally depth-modulated surfaces, and the perceived depth and shearing magnitudes of square wave depth-modulated surfaces. These data were compared with the predictions of the models in order to determine whether the visual system complies, either qualitatively or quantitatively, with the geometry of the situation. Perceived rotation was measured using the paddle matching procedure described above. Depth matching was carried out using a pair of digital callipers, which observers adjusted to match the perceived peak-to-trough amplitude of the surface.

Experiment 11 examined the perceived magnitudes of depth and rotation as a function of relative motion amplitude for surfaces of corrugation frequency 0.05, 0.1, 0.2 or 0.4 cpd. Increases in corrugation frequency led to increases in perceived rotation magnitude (as found in Experiments 6 and 7) which were accompanied by decreases in the perceived depth magnitude. This result is qualitatively consistent with a depth-rotation trade-off relationship. Comparison of the data with the model showed that the combined amounts of depth and rotation were significantly greater than predicted, however. In some cases the relative

motion amplitude theoretically required to account for the observed magnitudes of depth and rotation was almost three times the physical amplitude of the surface.

The perceived depth of surfaces has previously been found to scale with the projected width of the structure. For example, Caudek and Proffitt (1993) found that a horizontally oriented dihedral angle was perceived as deeper when the stimulus was wider. Rotation measurements were not taken in that study, however, and it was noted (Section 5.2.4) that the increase in motion gradient accompanying the decrease in wedge width would be expected to result in an increase in the magnitude of perceived rotation. If the visual system adheres, at least qualitatively, to a depth/rotation trade-off relationship then the decreased magnitudes of depth perceived in the smaller stimuli are predicted.

Experiment 12 determined the magnitudes of perceived depth and rotation for stimuli ranging in size from 2.5° to 20°. Consistent with the results of Experiment 9, perceived rotation magnitude increased with decreasing stimulus size. Concomitant decreases in the magnitude of perceived depth were also observed. Once more, these results were interpreted as qualitative support for the geometric trade-off model. Again, however, the relative motion amplitudes theoretically required to support the observed sizes of the two attributes were substantially larger than the physical amplitudes of the surfaces. The extent of the over-estimation was slightly smaller in Experiment 12 compared with Experiment 11. It was hypothesised that this was because of the generally lower magnitudes of perceived depth in the Experiment 12.

Taken together, the results of Experiments 11 and 12 provide qualitative support for the geometric trade-off model. That perceived depth and rotation magnitudes were found to be inter-dependent is suggestive that the visual system is, to some extent, constrained by the geometry of the situation. The trade-off relationship which exists between these attributes illustrates the metameric qualities of the proximal stimulus, with the actual solution chosen by the visual system being dependent, at least in part, upon additional perspective and motion gradient information. The data of Experiments 11 and 12 suggest that observers have a consistent tendency to over-estimate the combined

extents of perceived depth and rotation relative to those predicted by the model. The results of a number of other studies (e.g. Ono et al., 1986; Ono and Steinbach, 1990; Sakuri and Ono, 1996), when interpreted with respect to the model, are consistent with the observed over-estimations. Three possible explanations of the over-estimation were considered. The first of these is an over-estimation of the 2-D relative dot motion — observed by both Ono and Steinbach (1990) and Sato (1989). Since 2-D image motion has been shown to influence the 3-D percept (e.g. Bradshaw et al, 1991; Eagle and Blake, 1995; Werkhoven and van Veen, 1995; Hogervorst and Eagle, 1998) it is possible that biases in the initial motion analyses could lead to the observed over-estimations of depth and rotation magnitudes. A second possible explanation of the over-estimations of depth and rotation, namely that the relative motion was interpreted as both depth *and* rotation, was also suggested. This explanation involves the notion that when attention is directed towards the surface attribute currently being judged (depth, for example), the majority of the relative motion is interpreted as arising from that attribute. When attention is shifted to the other attribute (e.g. rotation) the amount of relative motion contributing to the magnitude of that percept may be larger than that theoretically 'used up' by the depth percept. This cannot be the entire story, however, since Experiments 11 and 12 showed that the combined magnitudes of depth and rotation were frequently greater than twice the physical amplitude of the surface. A combination of an over-estimation of the extent of 2-D relative motion and a propensity of the visual system to perceive a given amount of relative motion as both depth *and* rotation, would account for the observed results. The third possibility considered was that biases arose due to the measuring techniques. It is entirely possible that observers consistently over-represented the perceived magnitudes of rotation and depth in motion parallax surfaces, and there is no independent test of the extent to which this may be true. It is perhaps likely that all three of the above factors played a role in the observed over-estimations.

Experiment 13 investigated the perceived magnitudes of depth and shearing in square wave parallax surfaces, in order to determine whether a trade-off relationship exists between these attributes. A greater extent of depth was perceived in the surface of lower corrugation frequency. However, in contrast with the results of Experiment 11, there was no depth over-estimation for this surface. The extents of perceived shearing

were slightly higher in the surface of higher corrugation frequency. These data are consistent with a trade-off between perceived depth and shearing magnitudes. The combined magnitudes of these attributes were generally greater than those predicted by the geometric model, although the extent of this over-estimation was not as large as in Experiments 11 and 12. The relative motion amplitudes theoretically required to support the observed magnitudes of the attributes were generally less than twice the physical amplitude. In keeping with reports of over-estimation of 2-D dot motion (Sato, 1989; Ono and Steinbach, 1990), the perceived extent of shearing alone was greater than the relative motion of the display. This provides some support for the notion that an over-estimation of the extent of 2-D relative motion may contribute to the over-estimation of the combined magnitudes of depth and rotation in sinusoidal motion parallax surfaces.

Experiment 14 measured the extents of perceived depth and shearing for stimulus sizes of 2.5° and 20°. Perceived depth was greater in the larger stimulus. Increases in relative motion amplitude beyond 20 arc min equivalent disparity did not lead to further increases in the perceived depth magnitude and indeed for one subject led to an increasingly two-dimensional interpretation of the display. The perceived magnitude of shearing was greater in the larger surface, providing evidence against a trade-off between depth and shearing under these conditions. Once again, the shearing magnitudes were generally greater than the extents of 2-D relative dot motion. The combined extents of depth and shearing were greater than predicted by the geometric model, but again were significantly smaller than those of the comparable surfaces of Experiment 12.

A comparison of the results for sine and square wave surfaces illustrates the importance of the type of motion gradient in the determination of the percept obtained. Depth was perceived in sinusoidally modulated surfaces of every corrugation frequency and size examined. In contrast, the square wave pattern of relative motion was frequently perceived to contain less depth than predicted, with the visual system demonstrating a propensity to interpret the motion as 2-D shearing rather than 3-D structure. That a depth/shearing trade-off was evident when corrugation frequency was manipulated, but was *not* observed when stimulus size was varied, suggests that the 3-D percepts evoked by a square wave gradient are less predictable than those evoked by smoother — for

instance sinusoidal — motion gradients. It would appear that the task of disentangling depth and shearing magnitudes is a more difficult one than that of disentangling depth and rotation magnitudes. This notion is supported by subjective reports, noted at the time of the experiments, concerning the difficulty of the task.

Experiment 15 investigated the perceived magnitudes of depth and rotation as a function of viewing distance. It has been demonstrated that depth constancy in motion parallax surfaces exists up to distances of approximately 1 m (Ono et al., 1986; Rivest et al., 1989; see also Section 1.3.6). At greater viewing distances, the magnitude of perceived depth is substantially less than simulated. This contrasts with the case in stereo surfaces: in a well-lit, structured environment constancy can be close to 100 per cent (Glennerster et al., 1996). Durgin et al. (1995) showed that observers could make accurate shape judgements of real objects at viewing distances of up to 3 m. A possible reason for the significantly worse performance with parallax stimuli is the apparent rotation of the structures. Ono et al. (1986) found that an increased number of observers perceived motion parallax surfaces to rotate as the viewing distance increased. If rotation is perceived at the expense of depth then it would be predicted that the amount of perceived depth would fall off with increasing viewing distance. Experiment 15 investigated the perceived magnitudes of depth and rotation in proximally identical stimuli. It was found that for a given extent of proximal relative motion, the perceived extents of both depth and rotation increased with increasing distance. Consonant with the results of Ono et al. (1986) and Rivest et al. (1989), the magnitude of perceived depth was less than simulated at the distances of 114 and 228 cm. However, despite the fact that depth was increasingly under-estimated as viewing distance increased, the concomitant increases in the amount of perceived rotation were sufficient to over account for the relative motion 'unused' as depth. Thus as in the experiments at 57 cm, the combined extents of depth and rotation were greater than those predicted by the geometric model.

7.2.4 Chapter 6

The experiments of Chapter 5 suggest that a trade-off relationship exists between the perceived magnitudes of depth and rotation. The

concomitant variations in these attributes were produced through manipulations of corrugation frequency and stimulus size. The results of that chapter might lead one to conclude that any manipulation which altered the perceived extent of one attribute would necessarily lead to an alteration (in the opposite direction) of the magnitude of the other attribute. The experiments of Chapter 6 were designed to test whether a depth/rotation trade-off relationship occurs when a dynamic vertical perspective cue was introduced as a cue to rotation.

Experiment 16 examined perceived depth and rotation magnitudes when the rotation angle specified by the vertical perspective cue was ±15° and when it was 0° (i.e. when the cue was absent). Surfaces had a corrugation frequency of 0.05 - 0.4 cpd. An electronic match stimulus was used for the depth match, since pilot experiments using the digital callipers found that the data arising from these were relatively noisy under these conditions. Observers always perceived a greater extent of rotation when the perspective cue was present than when it was absent. This is an unsurprising but novel result, since there is no quantitative published work on the effects of dynamic vertical perspective in head-linked motion parallax displays. Increases in relative motion amplitude led to increases in perceived rotation, indicating that relative motion continues to play a role in perceived rotation even in the presence of the strong dynamic perspective cue. Despite the large effect of perspective upon perceived rotation magnitude (rotation matches were approximately 10° higher when the cue signalled a rotation of 15° than when it was absent), the magnitudes of depth were little affected. For the 0.2 and 0.4 cpd surfaces there was a small effect, such that less depth was perceived in the 15° condition than in the 0° condition.

Experiment 17 investigated perceived depth and rotation magnitudes for surfaces in which the dynamic vertical perspective cue indicated angles of 15° or 0°, for stimulus sizes of 5°, 10° and 20°. When the perspective cue was absent, the size 5° stimulus was perceived to rotate most and the size 20° stimulus was perceived to rotate least. When the cue signalled a rotation angle of 15° this trend was reversed, such that an increasing amount of rotation was perceived with increasing display size. These results are consistent with the increasing importance of vertical perspective in the determination of perceived rotation as stimulus size

increases. For the size 20° surface, perceived depth magnitudes were slightly smaller when the perspective cue signalled a 15° rotation than when it signalled a 0° rotation. For the smaller sizes there were generally no differences between the extents of perceived depth in the two situations.

Overall, the results of Experiments 16 and 17 suggest that variations in the perceived rotation magnitude do not necessarily go hand in hand with variations in perceived depth. Despite the substantial ability of the vertical perspective cue to determine perceived rotation under a variety of conditions, concomitant variations in depth were observed in only a subset of these conditions — namely those in which the stimulus size was relatively large and the corrugation frequency was relatively high. Even in these situations the variations in perceived depth were generally smaller than those observed in Experiments 11 and 12, when perceived rotation was varied via manipulations of corrugation frequency and stimulus size. The lack of a depth/rotation trade-off when perceived rotation magnitudes were manipulated using this cue must temper the notion that the visual system generally behaves in accord with the geometry, since in this case it did not.

Experiments 18 and 19 approached the depth/rotation relationship question using a shape judgement task. For five different vertical perspective levels, observers adjusted a triangle wave surface in order that the dihedral angle between the two sides of each peak (and each trough) appeared to be 90°. In Experiment 18 this was achieved through adjustment of the relative motion amplitude. When the rotation direction signalled by the vertical perspective cue was in the opposite direction to the head movement, the extent of relative motion required for the structure to contain perceived right-angles increased with increases in the simulated angle of rotation. This is consistent with a depth/rotation trade-off. When the vertical perspective cue indicated that the surface rotated in the same direction as the head motion, no clear pattern of depth results was obtained. This was likely to be due to the fact that observers did not perceive the surface as rotating in the simulated direction: rather they perceived it as a slightly non-rigid surface which rotated in the opposite direction to the head movement. For all surfaces, the relative motion amplitudes required for a perceived right-angled surface were lower than

predicted, indicating an over-estimation of perceived depth in keeping with the results of the experiments of Chapter 5.

Experiment 19 used the same shape judgement task but used surfaces of fixed relative motion amplitude and adjustable corrugation frequency. For the simulated convex rotations, the corrugation frequency required for a perceived right-angled surface increased with increasing simulated rotation angle. Once more, for the simulated concave rotations, non-rigidity and a convex rotation direction were perceived, and the pattern of results was equivocal. The corrugation frequencies required for the 'right-angled' surface were lower than predicted, again indicating an over-estimation of perceived depth magnitude. The fact that evidence suggestive of a depth/rotation trade-off was obtained in Experiments 18 and 19 but was obtained only under certain circumstances in Experiments 16 and 17 illustrates that the relationship is somewhat dependent upon the nature of the task. The fact remains, however, that the manipulation of either corrugation frequency or stimulus size results in substantial variations in the perceived depth magnitude whereas, under similar stimulus conditions and using an identical task, the manipulation of vertical perspective information does not.

Experiment 15 (Chapter 5) found that the perceived rotation magnitude of proximally identical stimuli increased with increasing viewing distance. A possible explanation of this is a decreasing reliance of the visual system upon the vertical perspective information (which specifies a stationary, frontal surface) as viewing distance increases, thus facilitating the perception of rotation at greater viewing distances. Experiment 20 tested this hypothesis by measuring perceived rotation at three viewing distances. The magnitudes of rotation specified by the vertical perspective cue were 0° and 15° at each distance. It was found that the effects of vertical perspective remained approximately constant over distance. The hypothesis that vertical perspective information is increasingly ignored as viewing distance increases was thus deemed unlikely. It was suggested that a factor possibly underlying the observed pattern of results is an increased propensity of the visual system to attribute vertical perspective changes to object rotation rather than to observer translation, with increasing viewing distance.

Experiment 21 examined the perceived depth order in surfaces in which the motion parallax was not accompanied by head translation. The observer remained stationary. In keeping with the literature, vertical perspective information was shown to affect perceived depth order. The novel finding of this experiment was that the proportion of response time spent perceiving a given depth order was dependent upon both the magnitude of rotation simulated by the vertical perspective cue and upon the stimulus size.

7.3 Discussion

These experiments provided a detailed analysis of the factors constraining the multiplicity of solutions which theoretically exist for a given pattern of relative motion. They illustrated that, wherever possible, the visual system interprets the motion as arising from a rigid, 3-D surface. When relative motion amplitudes are high or when motion gradients are steep, some of the relative motion is perceived as non-rigidity. A wide range of amplitudes gives rise to the percept of a non-rigid, 3-D surface, highlighting the willingness of the visual system to perceive depth (instead of, for example, 2-D shearing motion) even when this is necessarily accompanied by extreme non-rigidity.

At lower relative motion amplitudes, motion gradients and perspective information were both shown to be important factors in the interpretation of parallax motion. The relative weightings of the two cues in the determination of rotation thresholds and supra-threshold rotation magnitude were found to be different. Motion gradients were found to exert their influence at threshold rotation levels, whereas perspective information was the more important determinant of the perceived extent of supra-threshold rotation.

The perceived magnitudes of depth and rotation were found to co-vary when either corrugation frequency or stimulus size was manipulated. This was interpreted as qualitative evidence in support of the trade-off model of Chapter 5, and suggests that the visual system is constrained by a simple geometric analysis of the relative motion under these situations. The relationship between perceived magnitudes of depth and shearing in

square wave surfaces was not found to be so straightforward, and it is suggested that these attributes are less easily extracted than are the depth and rotation magnitudes of sinusoidal surfaces. This emphasises again the importance of the motion gradient in the percepts arising from motion parallax information.

Non-zero dynamic vertical perspective information was found to heavily influence the extent of perceived rotation when the simulated rotation direction was in the opposite direction to the head movement. The depth/rotation trade-off observed when perceived rotation was manipulated by varying corrugation frequency and stimulus size was not in evidence when perceived rotation was manipulated via the vertical perspective cue. This illustrates that the trade-off relationship is dependent upon the nature of the stimulus manipulation. The observation that sizeable variations in perceived rotation magnitude can occur *without* concomitant variations in the magnitude of perceived depth gives rise to the possibility that the apparent influence of corrugation frequency and stimulus size upon perceived depth may not necessarily be mediated through their effect upon perceived rotation magnitude. The observed trade-off, while being consistent with a causal relationship between the two attributes, does not rule out the possibility that the effects of corrugation frequency and stimulus size upon perceived depth and rotation are independent.

Surfaces for which the rotation simulated by the vertical perspective cue was in the same direction as the head movement were *not* perceived as rotating in the simulated direction. The vertical perspective transformations were instead interpreted as surface non-rigidity. This result emphasises the strength of the motion parallax information in the determination of perceived rotation direction. It illustrates that despite the existence of two geometrically legitimate rotation directions for a given pattern of relative motion, the visual system is biased towards perceiving rotation in the opposite direction to the head movement. This could, perhaps, be a consequence of an under-estimation of the extent of one's own motion.

In a situation in which motion parallax provided ambiguous information about depth order, vertical perspective information was

capable of influencing the proportion of time spent perceiving a given depth order. The importance of the simulated rotation *magnitude*, not merely its simulated direction, in the determination of perceived depth order was thus demonstrated.

It was noted at the outset that motion parallax is an inherently ambiguous cue. That the visual system does not have access to the infinite number of possible interpretations — or at least does not allow the observer to consciously perceive them — is apparent from the stable, compelling percepts arising from motion parallax information. This thesis has gone some way towards unravelling the factors upon which those percepts are consequent.

Appendix A

Calculation of the projected image of Point P following the rotation of the plane away from the frontal position.

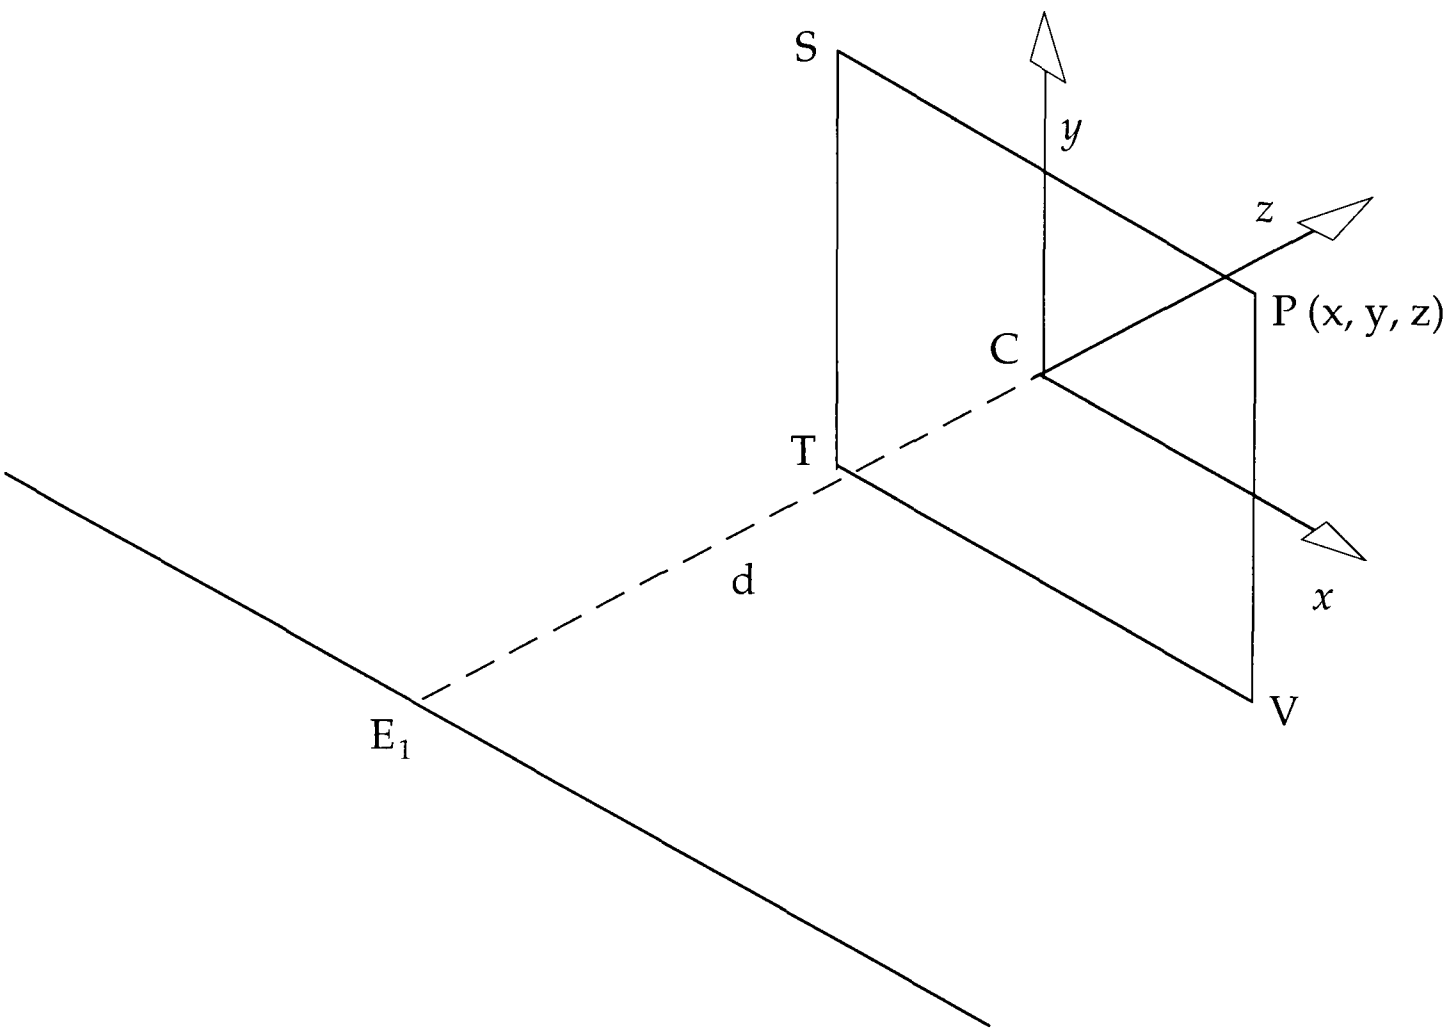


Fig. A.1: Plane SPVT in the frontal position, before rotation.

Fig. A.1 shows Plane SPVT in the frontal position, before rotation. C is in the centre of the plane. CE_1 is normal to the plane. CE_1 = d. The coordinates of Point P are (x, y, z).

As the eye moves from Point E_1 through a horizontal distance I to Point E_2, Plane SPVT rotates away from the frontal position, about a vertical axis through its centre C. The new positions of points S, P, V and T are S_2, P_2, V_2 and T_2.

Angle ϕ is made between Plane SPVT and Plane $S_2P_2V_2T_2$. Point P_2' is the projection of Point P_2 onto the frontal (projection) plane through E_2. Point S_2' is the projection of Point S_2 onto the frontal plane through E_2. The line $P_2'S_2'$ makes Angle α with the horizontal.

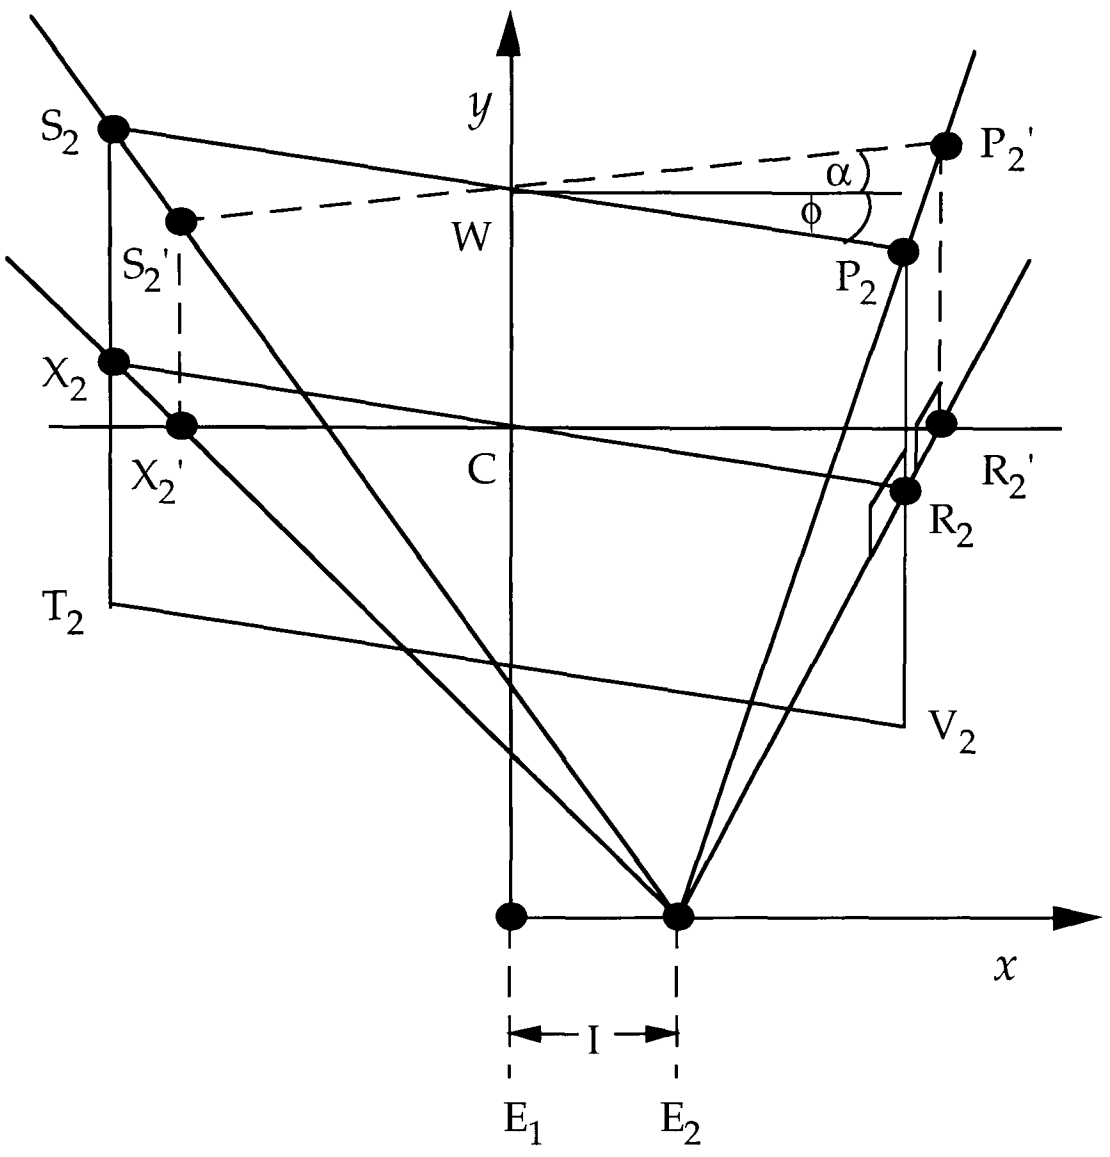


Fig. A.2: Plane $S_2P_2V_2T_2$, at a slant of ϕ with respect to the frontal plane.

Require the coordinates of P_2'.

To obtain the vertical coordinate of P_2':

In Fig. A.2:
Similar triangles $E_2P_2R_2$ and $E_2P_2'R_2'$:
$P_2'R_2'/P_2R_2 = E_2R_2'/E_2R_2 \therefore P_2'R_2' = (E_2R_2' * P_2R_2)/E_2R_2$.

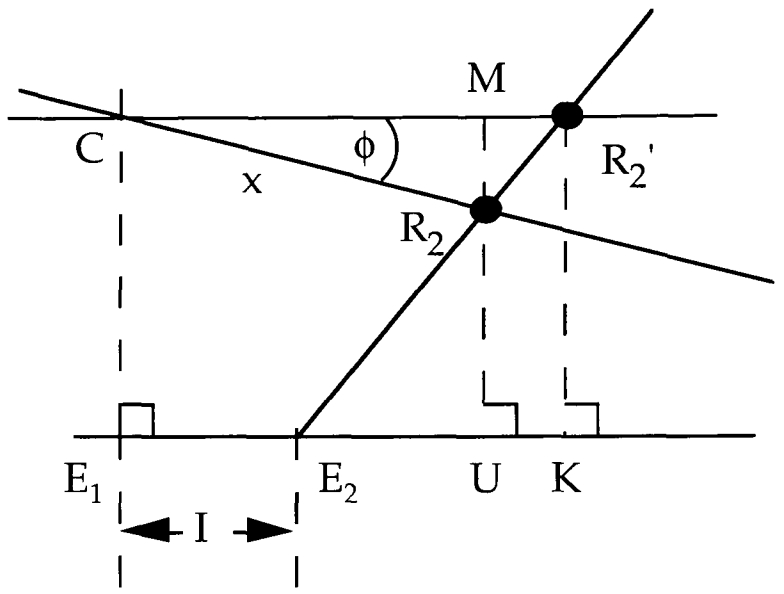


Fig. A.3: Plan view of the horizontal plane containing C.

In Fig. A.3:

$E_2R_2 = \sqrt{(E_2U^2 + R_2U^2)}$

$= \sqrt{[(x\cos\phi - I)^2 + (d - x\sin\phi)^2]}$

$E_2R_2' = E_2R_2 + R_2R_2'$

Similar triangles E_2R_2U and MR_2R_2':

$R_2R_2'/E_2R_2 = MR_2/R_2U \therefore R_2R_2' = (MR_2 * E_2R_2)/R_2U.$

$\therefore P_2'R_2' = \{[E_2R_2 + (MR_2 * E_2R_2)/R_2U] * P_2R_2\}/E_2R_2$

$= \{[(E_2R_2 * R_2U) + (MR_2 * E_2R_2)] * P_2R_2\}/(E_2R_2 * R_2U)$

$= [(R_2U + MR_2) * P_2R_2]/R_2U$

$= \{[(d - x\sin\phi) + x\sin\phi] * y\}/(d - x\sin\phi)$

$= dy/(d - x\sin\phi).$

To obtain the horizontal coordinate of P_2':

Require CR_2'.

$CR_2' = CM + MR_2'.$

Similar triangles MR_2R_2' and E_2R_2U:

$MR_2'/E_2U = MR_2/R_2U \therefore MR_2' = (MR_2 * E_2U)/R_2U$

$\therefore CR_2' = CM + [(E_2U * MR_2)/R_2U]$

$= x\cos\phi + \{[(x\cos\phi - I) * x\sin\phi]/(d - x\sin\phi)\}$

$= [x\cos\phi(d - x\sin\phi) + x\sin\phi(x\cos\phi - I)]/(d - x\sin\phi)$

$= (dx\cos\phi - Ix\sin\phi)/(d - x\sin\phi)$

$= x(d\cos\phi - I\sin\phi)/(d - x\sin\phi)$

The coordinates of P_2' are thus
$\{[x(d\cos\phi - I\sin\phi)/(d - x\sin\phi)], [dy/(d - x\sin\phi)]\}$.

We require RZ, such that Angle PWZ = α (see Fig. A.4). Recall that the electronically simulated rotations of the stimuli used in this thesis do *not* take into account the horizontal width changes which accompany real-life rotations (see Section 2.1.4).

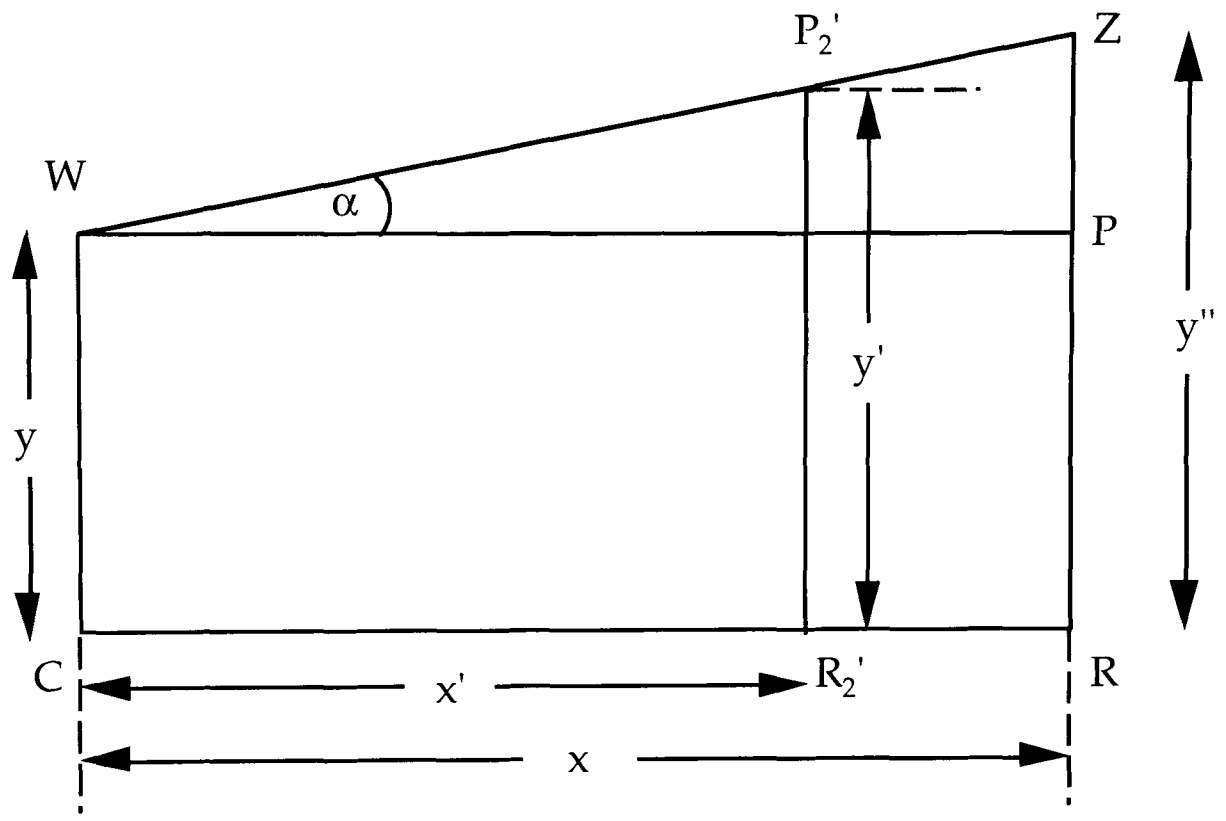


Fig. A.4: $WCP_2'R_2'$ is the frontal projection of the upper right-hand quadrant of rotated Plane $S_2P_2V_2T_2$. XPRC is the upper right-hand quadrant of frontal Plane SPVT.

$RZ = RP + PZ.$

$y'' = y + x\tan\alpha$

Require $\tan\alpha$.

In Fig. A.4,

$$\begin{aligned}\tan\alpha &= (y' - y)/x' \\ &= \{[dy/(d - x\sin\phi)] - y\} \,/\, \{[x(d\cos\phi - I\sin\phi)]/(d - x\sin\phi)\} \\ &= \{[dy - y(d - x\sin\phi)]/(d - x\sin\phi)\} \,/\{[x(d\cos\phi - I\sin\phi)]/(d - x\sin\phi)\} \\ &= [dy - y(d - x\sin\phi)]/[x(d\cos\phi - I\sin\phi)] \\ &= xy\sin\phi \,/\, [x(d\cos\phi - I\sin\phi)] \\ &= y\sin\phi \,/\, (d\cos\phi - I\sin\phi).\end{aligned}$$

Thus $y'' = y + [(xysin\phi)/(dcos\phi - Isin\phi)]$.

The voltage required to simulate a rotation angle of 15°, for a plane 18 cm x 18 cm, was calculated as follows:

Given that

$x = 9$ cm
$y = 9$ cm
$d = 57.1$ cm
$I = 6.5$ cm,

$$y'' = y + [(xysin\phi)/(dcos\phi - Isin\phi)]$$
$$= 9 + [(9 * 9sin15°)/(57.1cos15° - 6.5sin15°)]$$
$$= 9.39 \text{ cm}$$

$\therefore$ total height of trapezium's RHS
$= 9.39 * 2$
$= 18.78$ cm
$= 187.8$ mm.

The voltage required to produce a trapezium with a RHS of 187.8 mm is found by solving the following equation to find x (see Fig. 2.4):

$$y = 180.03 + 4.35x$$
$$\therefore 187.8 = 180.03 + 4.35x$$
$$\therefore x = (187.8 - 180.03)/4.35$$
$$= 1.79 \text{ volts.}$$

Note that since the horizontal width changes which would occur with rotation of a real-world stimulus are not simulated, the centre of a rotated plane which projects to the displayed trapezium does *not* lie on C (thus the use of the word 'quadrant' in the first instance in the caption of Fig. A.4 is not strictly correct).

Appendix B

Equations of the power functions which were fit to the data of Experiment 1b (Figs. 3.17 - 3.19).

Subject SAS

Unclipped	: y = 0.88633 * x^(0.18754)	R = 1
50%-clipped	: y = 1.1155 * x^(0.30103)	R = 1
25%-clipped	: y = 0.96721 * x^(-0.024197)	R = 1

Subject AMJ

Unclipped	: y = 2.2782 * x^(0.047842)	R = 1
50%-clipped	: y = 2.2901 * x^(-0.017465)	R = 1
25%-clipped	: y = 2.1983 * x^(-0.027776)	R = 1

Subject SJP

Unclipped	: y = 1.2362 * x^(-0.028105)	R = 1
50%-clipped	: y = 1.3313 * x^(-0.1772)	R = 1
25%-clipped	: y = 1.2303 * x^(-0.20533)	R = 1

Appendix C

To test the significance of the differences between Transition Points of different waveforms at each of the six corrugation frequencies.

$HSD = q_{(0.05)}(p, v) * \sqrt{(MS_{error}/n)}$

Where
n = number of observations in each mean
p = number of means
v = p(n-1).

In the tables below, the Transition Points of each waveform are shown in plain text. Italicised text represents differences between Transition Points. Significant (at the 0.05 level) differences are marked with an asterisk.

Subject SAS

n = 3
p = 4
v = 8

$q_{(0.05)}(4,8) = 4.53$

$\therefore HSD = 4.53 * \sqrt{(0.05/3)}$
$= 0.585$

0.05 cpd

Waveform		**Triangle**	**Sine**	**Square**	**Ramp**
	T. Point	*7.677*	8.779	16.425	16.716
Triangle	7.677	*0.000*			
Sine	8.779	*1.102**	*0.000*		
Square	16.425	*8.748**	*7.646**	*0.000*	
Ramp	16.716	*9.039**	*7.937**	*0.291*	*0.000*

0.1 cpd

Waveform		Triangle	Sine	Ramp	Square
	T. Point	4.557	4.573	8.353	11.286
Triangle	4.557	*0.000*			
Sine	4.573	*0.016*	*0.000*		
Ramp	8.353	*3.796**	*3.780**	*0.000*	
Square	11.286	*3.271**	*6.713**	*2.933**	*0.000*

0.2 cpd

Waveform		Triangle	Sine	Ramp	Square
	T. Point	2.688	3.253	6.610	7.032
Triangle	2.688	*0.000*			
Sine	3.253	*0.565*	*0.000*		
Ramp	6.610	*3.922**	*3.357**	*0.000*	
Square	7.032	*4.344**	*3.379**	*0.422*	*0.000*

0.4 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	1.136	1.261	3.524	7.797
Sine	1.136	*0.000*			
Triangle	1.261	*0.125*	*0.000*		
Ramp	3.524	*2.388**	*2.263**	*0.000*	
Square	7.797	*6.661**	*6.536**	*4.273**	*0.000*

0.8 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	0.852	1.073	3.145	4.829
Sine	0.852	*0.000*			
Triangle	1.073	*0.221*	*0.000*		
Ramp	3.145	*2.293**	*2.072**	*0.000*	
Square	4.829	*3.977**	*3.756**	*1.684**	*0.000*

1.6 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	0.944	1.475	3.432	4.604
Sine	0.944	*0.000*			
Triangle	1.475	*0.531*	*0.000*		
Ramp	3.432	*2.488**	*1.957**	*0.000*	
Square	4.604	*3.660**	*3.129**	*1.172**	*0.000*

Subject AMJ

n = 3

p = 4

v = 8

$q_{(0.05)}$ (4,8) = 4.53

$\therefore$ HSD = 4.53 * $\sqrt{(0.06/3)}$

= 0.641

0.05 cpd

Waveform		**Ramp**	**Sine**	**Square**	**Triangle**
	T. Point	8.040	11.174	12.171	12.246
Ramp	8.040	*0.000*			
Sine	11.174	*3.134**	*0.000*		
Square	12.171	*4.131**	*0.997**	*0.000*	
Triangle	12.246	*4.206**	*1.072**	*0.075*	*0.000*

0.1 cpd

Waveform		**Triangle**	**Sine**	**Ramp**	**Square**
	T. Point	6.927	7.266	7.654	10.258
Triangle	6.927	*0.000*			
Sine	7.266	*0.339*	*0.000*		
Ramp	7.654	*0.727**	*0.388*	*0.000*	
Square	10.258	*3.331**	*2.992**	*2.604**	*0.000*

0.2 cpd

Waveform		**Sine**	**Ramp**	**Triangle**	**Square**
	T. Point	4.542	5.120	5.305	8.377
Sine	4.542	*0.000*			
Ramp	5.120	*0.578*	*0.000*		
Triangle	5.305	*0.763**	*0.185*	*0.000*	
Square	8.377	*3.835**	*3.257**	*3.072**	*0.000*

0.4 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	2.388	2.940	3.398	7.515
Sine	2.388	*0.000*			
Triangle	2.940	*0.552*	*0.000*		
Ramp	3.398	*3.950**	*0.458*	*0.000*	
Square	7.515	*5.127**	*4.575**	*4.117**	*0.000*

0.8 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	2.240	2.925	2.935	6.802
Sine	2.240	*0.000*			
Triangle	2.925	*0.685**	*0.000*		
Ramp	2.935	*0.695**	*0.010*	*0.000*	
Square	6.802	*4.562**	*3.877**	*3.867**	*0.000*

1.6 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	2.222	2.535	2.804	5.667
Sine	0.222	*0.000*			
Triangle	2.535	*0.313*	*0.000*		
Ramp	2.804	*0.582*	*0.269*	*0.000*	
Square	5.667	*2.445**	*3.132**	*2.863**	*0.000*

<u>Subject SJP</u>

$n = 3$

$p = 4$

$v = 8$

$q_{(0.05)}(4,8) = 4.53$

$\therefore$ HSD $= 4.53 * \sqrt{(0.35/3)}$

$= 1.547$

0.05 cpd

Waveform		Triangle	Sine	Square	Ramp
	T. Point	9.322	9.824	22.508	23.026
Triangle	9.322	*0.000*			
Sine	9.824	*0.502*	*0.000*		
Square	22.508	*13.186**	*12.684**	*0.000*	
Ramp	23.026	*13.704**	*13.202**	*0.518*	*0.000*

0.1 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	6.333	6.428	9.853	17.691
Sine	6.333	*0.000*			
Triangle	6.428	*0.095*	*0.000*		
Ramp	9.853	*3.520**	*3.425**	*0.000*	
Square	17.691	*11.358**	*11.263**	*8.108**	*0.000*

0.2 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	3.716	4.037	5.867	11.268
Sine	3.716	*0.000*			
Triangle	4.037	*0.321*	*0.000*		
Ramp	5.867	*2.151**	*1.830**	*0.000*	
Square	11.268	*7.552**	*7.231**	*5.401**	*0.000*

0.4 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	1.843	2.396	4.705	9.395
Sine	1.843	*0.000*			
Triangle	2.396	*0.553*	*0.000*		
Ramp	4.705	*2.862**	*2.309**	*0.000*	
Square	9.395	*7.552**	*6.999**	*4.690**	*0.000*

0.8 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	1.234	1.971	3.588	9.234
Sine	1.234	*0.000*			
Triangle	1.971	*0.737*	*0.000*		
Ramp	3.588	*2.354**	*1.617**	*0.000*	
Square	9.234	*8.000**	*7.263**	*5.646**	*0.000*

1.6 cpd

Waveform		Sine	Triangle	Ramp	Square
	T. Point	1.204	1.969	2.922	8.048
Sine	1.204	*0.000*			
Triangle	1.969	*0.765*	*0.000*		
Ramp	2.922	*1.718**	*0.953*	*0.000*	
Square	8.048	*6.844**	*6.079**	*5.126**	*0.000*

Appendix D

Pilot data: depth matching using electronic callipers (Experiment 16), for Subject SAS. The four graphs show Depth Matches for surfaces of corrugation frequency 0.05, 0.1, 0.2 and 0.4 cpd respectively.

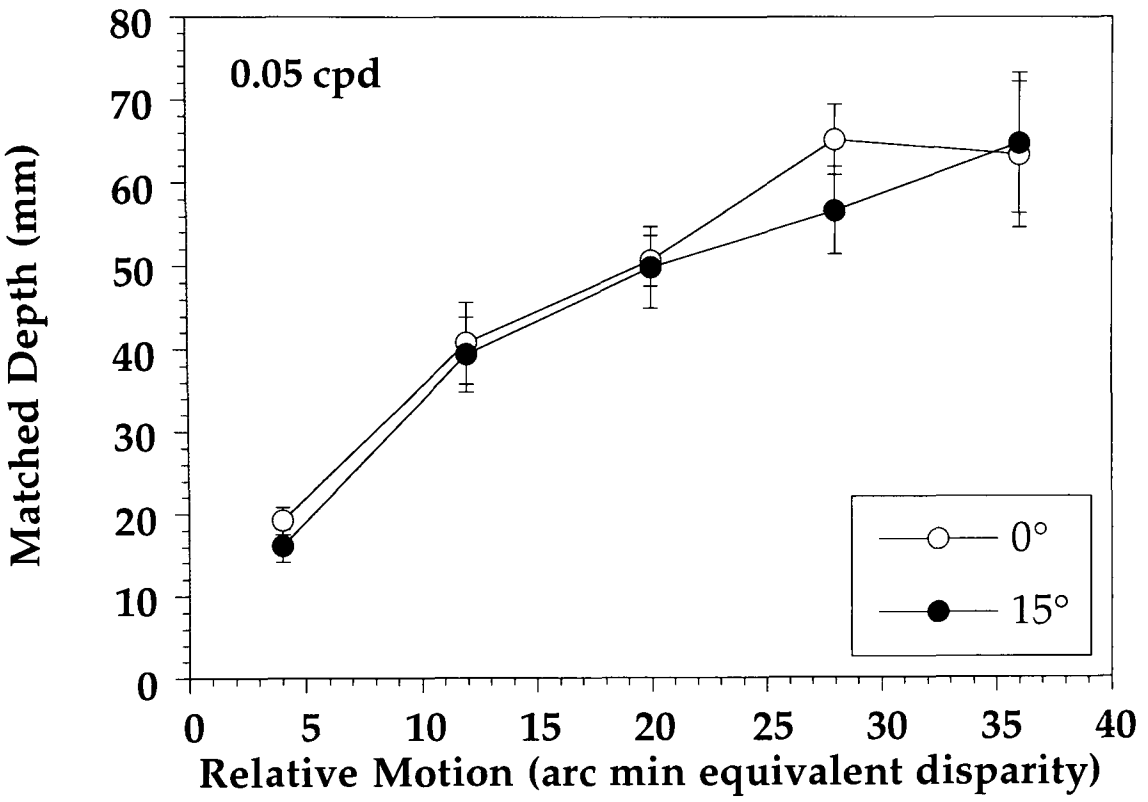


Fig. C.1: Matched Depth for the 0.05 cpd surface, for Subject SAS. The open symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 15°.

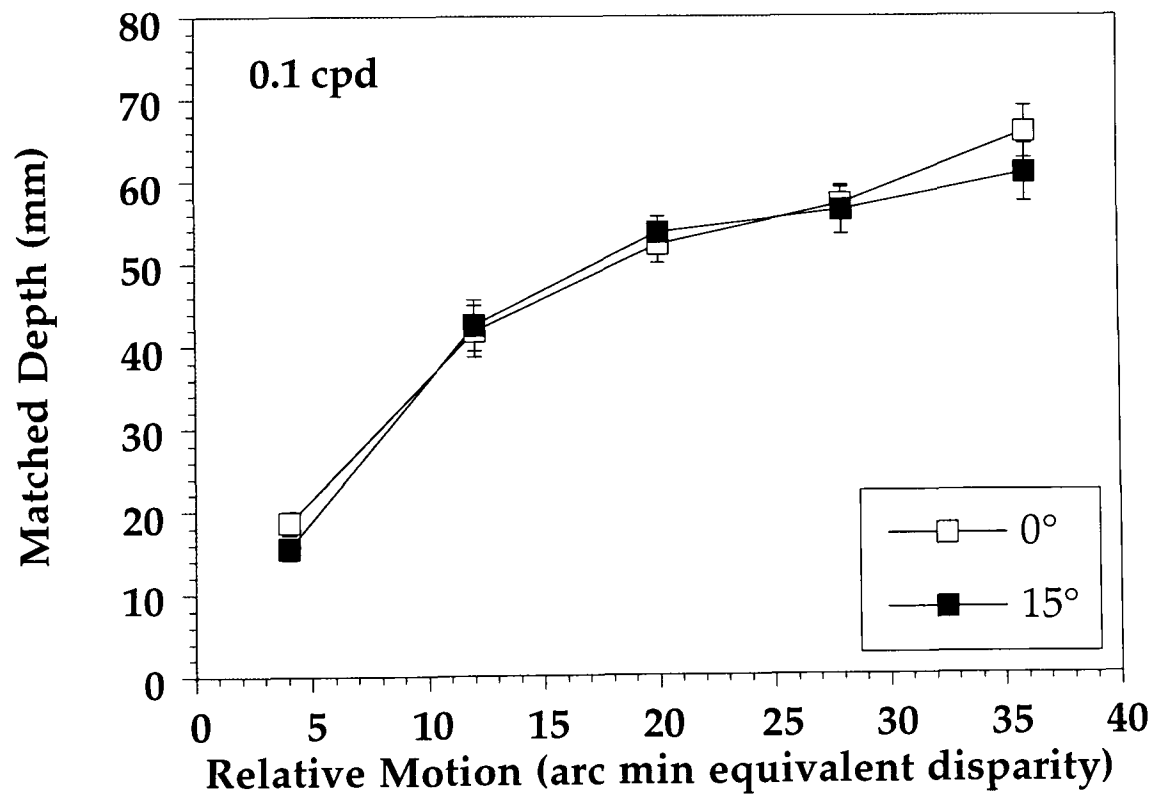


Fig. C.2: Matched Depth for the 0.1 cpd surface, for Subject SAS. The open symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 15°.

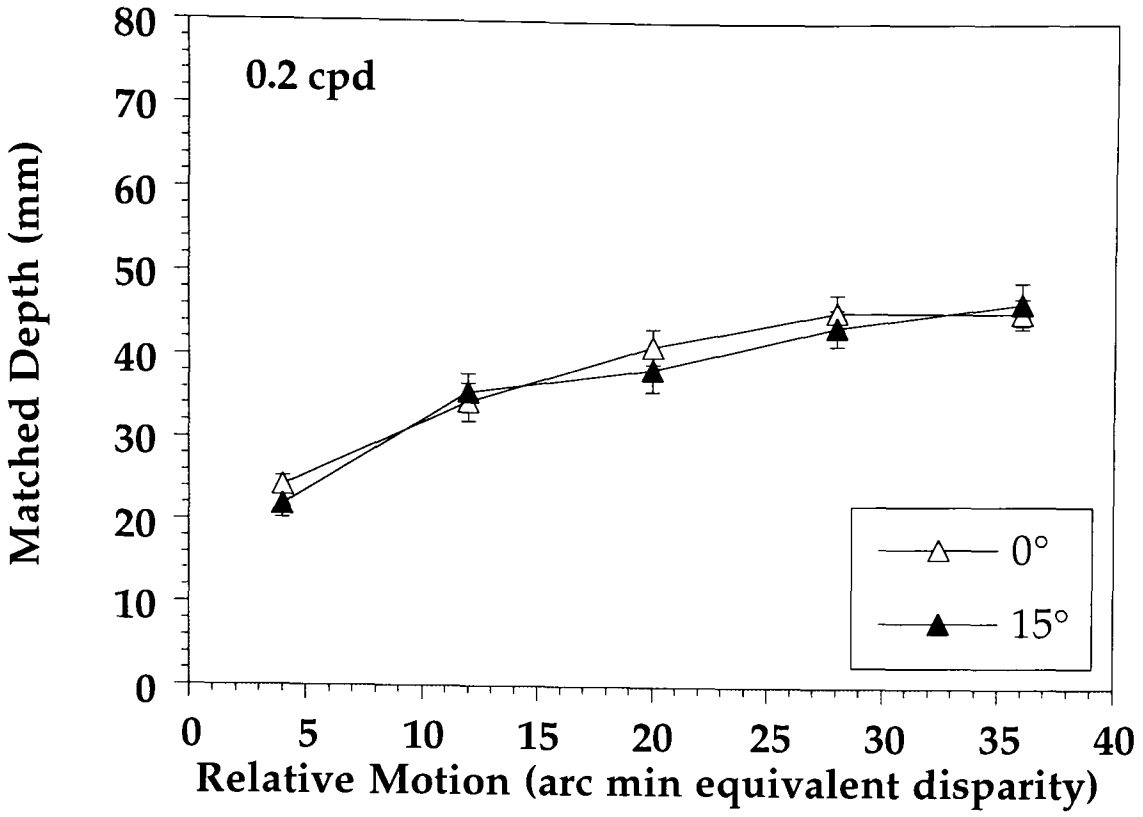


Fig. C.3: Matched Depth for the 0.2 cpd surface, for Subject SAS. The open symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 15°.

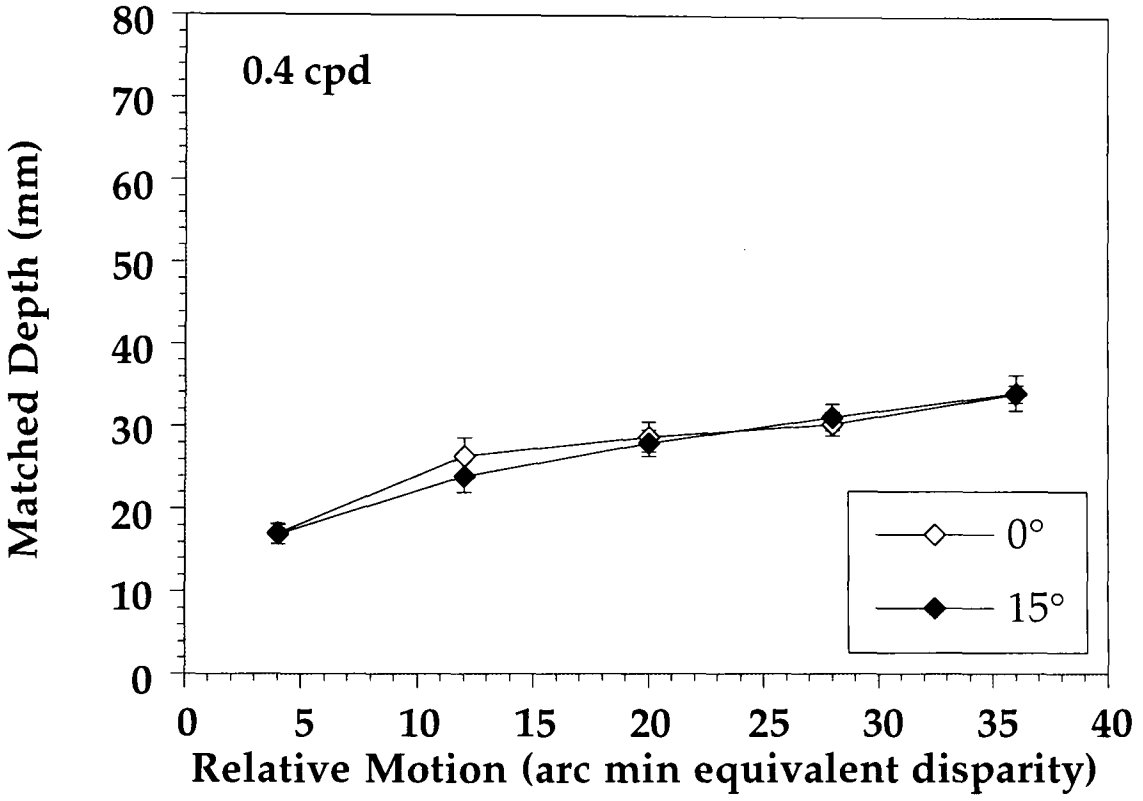


Fig. C.4: Matched Depth for the 0.4 cpd surface, for Subject SAS. The open symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 0°. The closed symbols show matches made in the condition in which vertical perspective simulated a rotation angle of 15°.

References

Adelson, E. H. and Movshon, J.A. (1982). Phenomenal coherence of moving visual patterns. *Nature*, **300**, 523-525.

Ames, A. (1951). Visual perception and the rotating trapezoidal window. *Psychological Monographs*, **67** (7, Whole No. 324).

Andersen, G.J., and Braunstein, M.L. (1983). Dynamic occlusion in the perception of rotation in depth. *Perception & Psychophysics*, **34**, 356-362.

Andersen, G.J. and Braunstein, M.L. (1985). Induced self-motion in central vision. *Journal of Experimental Psychology: Human Perception and Performance*, **11**, 122-132.

Anstis, S.M., Howard, I.P. and Rogers, B.J. (1978). A Craik-Cornsweet illusion for visual depth. *Vision Research*, **18**, 213-217.

Arditi, A. (1982). The dependence of the induced effect on orientation and a hypothesis concerning disparity computations in general. *Vision Research*, **22**, 247-256.

Ball, W. and Tronick, E. (1971). Infant responses to impending collision: Optical and real. *Science*, **171**, 818-820.

Barnard, S.T. and Thompson, W.B. (1980). Disparity analysis of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAM1-2**, 333-340. Cited in Ono, Rivest and Ono (1986).

Bassili, J.N. (1978). Facial motion in the perception of faces and of emotional expression. *Journal of Experimental Psychology: Human Perception and Performance*, **4**, 373-379.

Berbaum, K., Bever, T. and Chung, C.S. (1983). Light source position in the perception of object shape. *Perception*, **12**, 411-416.

Berbaum, K., Bever, T. and Chung, C.S. (1984). Extending the perception of shape from known to unknown shading. *Perception*, **13**, 479-488.

Berbaum, K., Tharp, D. and Mroczek, K. (1983). Depth perception of surface in pictures: Looking for conventions of depiction in Pandora's box. *Perception*, **12**, 5-20.

Berkeley, G. (1709). *An essay towards a new theory of vision.* Dublin, Jeremy Pepyat. Reprinted 1922. New York, Dutton. Cited in Howard and Rogers (1995).

Beverley, K.I. and Regan, D. (1980). Separable aftereffects of changing-size and motion in depth: different neural mechanisms? *Vision Research*, **19**, 727-732.

Boring, E.G. (1942). *Sensation and perception in the history of psychology*. Appleton-Century-Crofts, New York.

Börjesson, E. (1971). Properties of changing patterns evoking visually perceived oscillation. *Perception & Psychophysics*, **9**, 303-308.

Bourdon, B. (1902). *La perception visuelle de l'espace*. Paris: Schleicher, 1906. Cited in Gibson, Gibson, Smith and Flock (1959).

Bower, T.G.R., Brougton, J.M. and Moore, M.K. (1970). Infant responses to approaching objects. An indicator of response to distal variables. *Perception & Psychophysics*, **9**, 193-196.

Bradley, D.C., Qian, N. and Anderson, R.A. (1995). Integration of motion and stereopsis in middle temporal cortical area of macaques. *Nature*, **373**, 609-611.

Bradshaw, M., Glennerster, A. and Rogers, B.J. (1996). The effect of display size on disparity scaling from differential perspective and vergence cues. *Vision Research*, **36**, 1255-1264.

Bradshaw, M.F.B. and Rogers, B.J. (1993). Elevation of depth thresholds following within- and between-cue adaptation of stereo and motion. *IOVS*, **34**, 1036.

Bradshaw, M.F. and Rogers, B.J. (1994). Cited in Rogers and Howard (1995).

Bradshaw, M.F.B. and Rogers, B.J. (1996). The interaction of binocular disparity and motion parallax in the computation of depth. *Vision Research*, **36**, 3457-3468.

Bradshaw, M.F., Rogers, B.J.R., Ono, H. and Ohmi, M. (1991). Thresholds for perceiving (i) 3-D structure from motion parallax and (ii) 2-D relative motion as a function of dot lifetime. *IOVS*, **32**, 811.

Braunstein, M.L. (1962). Depth perception in rotating dot patterns: effects of numerosity and perspective. *Journal of Experimental Psychology*, **64**, 415-420.

Braunstein, M.L. (1966). Sensitivity of the observer to transformations of the visual field. *Journal of Experimental Psychology*, **72**, 683-687.

Braunstein, M.L. (1968). Motion and texture as sources of slant information. *Journal of Experimental Psychology*, **73**, 247-253.

Braunstein, M.L. (1971). Perception of rotation in figures with rectangular and trapezoidal features. *Journal of Experimental Psychology*, **91**, 25-29.

Braunstein, M.L. (1972). Perception of rotation in depth: a process model. *Psychological Review*, **79**, 510-524.

Braunstein, M.L. (1976). *Depth perception through motion*. Academic Press, New York.

Braunstein, M.L. (1977a). Minimal conditions for the perception of rotary motion. *Scandinavian Journal of Psychology*, **18**, 216-223.

Braunstein, M.L. (1977b). Perceived direction of rotation of simulated three-dimensional patterns. *Perception & Psychophysics*, **21**, 553-557.

Braunstein, M.L (1994). Structure from motion. In A.T. Smith and R.J. Snowden (eds.), *Visual detection of motion* (pp. 367-393). New York: Acacemic Press.

Braunstein, M.L. and Andersen, G.J. (1981). Perceived relative depth from velocity gradients: a model. *Acta Psychologica*, **48**, 195-201.

Braunstein, M.L. and Andersen, G.J. (1981). Velocity gradients and relative depth perception. *Perception & Psychophysics*, **29**, 145-155.

Braunstein, M.L. and Andersen, G.J. (1984). A counter example to the rigidity assumption in the perception of structure from motion. *Perception*, **13**, 213-217.

Braunstein, M.L. and Andersen, G.J. (1984). Shape and depth perception from parallel projections of three-dimensional motion. *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 749-760.

Braunstein, M.L, Andersen, G.F and Riefer, D.M. (1982). The use of occlusion to resolve ambiguity in parallel projections. *Perception & Psychophysics*, **31**, 261-267.

Braunstein, M.L., Andersen, G.J., Rouse, M.W. and Tittle, J.S. (1986). Recovering viewer-centred depth from disparity, occlusion, and velocity gradients. *Perception & Psychophysics*, 40, 216-224.

Braunstein, M.L., Hoffman, D.D. and Pollick, F.E. (1990). Discriminating rigid from non-rigid motion: minimum points and views. *Perception & Psychophyics*, **47**, 205-214.

Braunstein, M.L., Hoffman, D.D., Shapiro, L.R., Andersen, G.J. and Bennett, B.M. (1987). Minimum points and views for the recovery of three-dimensional structure. *Journal of Experimental Psychology: Human Perception and Performance*, **13**, 335-343.

Braunstein, M.L., Liter, J.C. and Tittle, J.S. (1993). Recovering three-dimensional shape from perspective translations and orthographic rotations. *Journal of Experimental Psychology: Human Perception and Performance*, **19**, 598-614.

Braunstein, M.L. and Payne, J.W. (1968). Perspective and the rotating trapezoid. *Journal of the Optical Society of America*, **58**, 399-403.

Braunstein, M.L. and Stern, K.R. (1980). Static and dynamic factors in the perception of rotary motion. *Perception & Psychophysics*, **27**, 313-320.

Braunstein, M.L. and Tittle, J.S. (1988). The observer-relative velocity field as the basis for effective motion parallax. *Journal of Experimental Psychology: Human Perception and Performance*, **14**, 582-590.

Braunstein, M.L. and Todd, J.T. (1990). On the distinction between artifacts and information. *Journal of Experimental Psychology: Human Perception and Performance*, **16**, 211-216.

Bülthoff, H., Fahle, M. and Wegmann, M. (1991). Perceived depth scales with disparity gradient. *Perception*, **20**, 145-53.

Bülthoff, H.H. and Mallot, H.A. (1988). Integration of depth modules: stereo and shading. *Journal of the Optical Society of America A*, **5**, No. 10, 1749-1758.

Burt, P. and Julesz, B. (1980). Modifications of the classical notion of Panum's fusional area. *Perception*, **9**, 671-682.

Cagenello, R. and Rogers, B.J. (1993). Anisotropies in the perception of stereoscopic surfaces: the role of orientation dispariy. *Vision Research*, **33**, 2189-2201.

Cambell, F.W. and Robson, J.G. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, **197**, 551-566.

Caudek, C. and Proffit, D.R. (1993). Depth perception in motion parallax and stereokinesis. *Journal of Experimental Psychology: Human Perception and Performance*, **19**, 32-47.

Clark, W.C., Smith, A.H. and Rabe, A. (1956). Retinal gradients of outline distortion and binocular disparity as stimuli for slant. *Candian Journal of Psychology*, **10**, 77-81.

Cobo–Lewis, A.B. and Yeh, Y.Y. (1994). Selectivity of cyclopean masking for the spatial frequency of binocular disparity modulation. *Vision Research*, **34**, 607-620.

Cords, R. (1913). Der Einfluss der parallaktischen Verscheibung auf die Tiefenwahrnehmung. *Klin. Mbl. Augenheilk.*, **51**, 421. Cited in Rogers (1993).

Cormack, R.H. (1984). Stereoscopic depth perception at far viewing distances. *Perception & Psychophysics*, **35**, 425-428.

Cornilleau-Pérès, V. and Droulez, J. (1989). Visual perception of surface curvature: psychophysics of curvature detection induced by motion parallax. *Perception & Psychophysics*, **46**, 351-364.

Cornilleau-Pérès, V. and Droulez, J. (1993). Stereo-motion cooperation and the use of motion disparity in the visual perception of 3D structure. *Perception & Psychophysics*, **54**, 223-239.

Cortese, J.M. and Andersen, G.J. (1991). Recovery of 3-D shape from deforming contours. *Perception*, **49**, 315-327.

Cumming, B.G., Johnston, E.B. and Parker, A.J. (1991). Vertical disparities and the perception of three-dimensional shape. *Nature*, **349**, 411-413.

Cutting, J.E. (1986). *Perception with an eye to motion*. MIT Press, Cambridge, MA.

Cutting, J.E. and Millard, R.T. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General*, **113**, 198-216.

Cutting, J.E., Proffitt, D.R. and Kozlowski, L.T. (1978). A biomechanical invariant for gait perception. *Journal of Experimental Psychology: Human Perception and Performance*, **4**, 357-372.

Dees, J.W. (1966). Accuracy of absoute visual distance and size estimation in space as a function of stereopsis and motion parallax. *Journal of Experimental Psychology*, **72**, 466-476.

Descartes, R. (1664). Traité de l'homme. In *Oeuvres de Descartes* volume XI, (ed. C. Adam and P. Tannery), 1909, p. 119-215. Cerf, Paris. Cited in Howard and Rogers (1995).

de Vries, S.C., Kappers, A.M.L. and Koenderink, J.J. (1993). Shape from stereo: A systematic approach using quadratic surfaces. *Perception & Psychophysics*, **53**, 71-80.

Durgin, F.H., Proffit, D.R., Reinke, K.S. and Olson, T.J. (1995). Comparing depth from motion with depth from binocular disparity. *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 679-699.

Dyson, F.W. (1929). Parallax. *Encyclopaedia Brittania*, **17**, 261-262.

Eagle, R.A. and Blake, A. (1994). 2-D limits on 3-D structure-from-motion tasks. *IOVS*, 1277.

Eagle, R.A. and Blake, A. (1995). Two-dimensional constraints on three-dimensional structure from motion tasks. *Vision Research*, **35**, 2927-2941.

Eagle, R.A. and Hogervorst, M. (1997). The recovery of structure from motion under perspective and orthographic projection. AVA, Surrey.

Eagle, R. A. and Hogervorst, M.A. (1997). The role of perspective information in structure-from-motion judgements. *IOVS*, 360.

Eby, D.W. (1992). The spatial and temporal characteristics of perciving 3-D structure from motion. *Perception and Psychophysics*, **51**, 163-78.

Eby, D.W. and Braunstein, M.L. (1993). The perceptual flattening of 3-D scenes enclosed by a frame. *IOVS*, **34**, 1184,

Eby, D.W. and Braunstein, M.L. (1995). The perceptual flattening of three-dimensional scenes enclosed by a frame. *Perception*, **24**, 981-993.

Epstein, W. and Mountford, D. (1963). Judgement of slant in response to an isolated gradient of stimulations. *Perceptual and Motor Skills*, **16**, 733-737.

Eriksson, E.S. (1973). Distance perception and the ambiguity of visual stimulation: a theoretical note. *Perception & Psychophysics*, **13**, 379-381.

Eriksson, E. Sture (1974). Movement parallax during locomotion. *Perception & Psychophysics*, **16**, 197-200.

Farber, J.M. and McConkie, A.B. (1979). Optical motions as information for unsigned depth. *Journal of Experimental Psychology: Human Perception and Performance*, **5**, 494-500.

Ferris, S.H. (1972). Motion parallax and absolute distance. *Journal of Experimental Psychology*, **95**, 258-63.

Fieandt, K. von and Gibson, J.J. (1959). The sensitivity of the eye to two kinds of continuous transformation of a shadow-pattern. *Journal of Experimental Psychology*, **57**, 344-347.

Finney, D.J. (1971). *Probit Analysis*. Cambridge: Cambridge University Press.

Fischer, G.J. (1956). Factors affecting estimations of depth with variations of the stereokinetic effect. *American Journal of Psychology*, **69**, 252-257.

Flock, H.R. (1964). Some conditions sufficient for accurate monocular perceptions of moving surface slants. *Journal of Experimental Psychology*, **95**, 258-263.

Foley, J.M. (1980). Binocular distance perception. *Psychological Review*, **87**, 411-434.

Foley, J.M. and Richards, W. (1972). Effects of voluntary eye movement and convergence on the binocular appreciation of depth. *Perception & Psychophysics*, **11**, 423-427.

Freeman, T.C.A., Harris, M.G. and Meese, T.S. (1996). On the relationship between deformation and perceived surface slant. *Vision Research*, **36**, 317-422.

Freeman, T.C.A., Harris, M.G. and Tyler, P.A. (1994). Human sensitivity to temporal proximity: the role of spatial and temporal speed gradients. *Perception & Psychophysics*, **55**, 689-699.

Ganis, G., Casco, C. and Roncat, S. (1993). Perceived rigidity and non-rigidity in the kinetic depth effect. *Perception*, **22**, 23-34.

Gescheider, G.A. (1997). *Psychophysics: The Fundementals* (third edition). Lawrence Erlbaum Associates.

Gibson, J.J. (1950). *The perception of the visual world*. Houghton, Mifflin, Boston, MA.

Gibson, J.J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin.

Gibson, J.J. and Bridgeman, B. (1987). The visual perception of surface texture in photographs. *Psychological Research*, **49**, 1-5.

Gibson, J.J. and Carel, W. (1952). Does motion perspective independently produce the impression of a receding surface? *Journal of Experimental Psychology*, **44**, 16-18.

Gibson, J.J. and Gibson, E.J. (1957). Continuous perspective transformations and the perception of rigid motion. *Journal of Experimental Psychology*, **54**, 129-138.

Gibson, E.J., Gibson, J.J., Smith, O.W. and Flock., H. (1959). Motion parallax as a determinant of perceived depth. *Journal of Experimental Psychology*, **58**, 40-51.

Gibson, J.J., Olum, P. and Rosenblatt, F. (1955). Parallax and perspective during aircraft landings. *American Journal of Psychology*, **68**, 373-385.

Gillam, B.J. (1968). Perception of slant when perspective and stereopsis conflict: experiments with aniseikonic lenses. *Journal of Experimental Psychology*, **78**, 299-305.

Gillam, B. and Lawergren, B. (1983). The induced effect, vertical disparity, and stereoscopic theory. *Perception & Psychophysics*, **34**, 121-130.

Gillam, B.J., Palmisano, S.A. and Rogers, B.J. (1998). Stereoscopic slant reversals occur for hoizontal disparity and overall disparity but not for vertical disparity. ECVP 1998.

Gillam, B. and Ryan, C (1992). Perspective, orientation disparity and anisotropy in stereoscopic slant perception. *Perception*, **21**, 427-439.

Glennerster, A. Rogers, B.J. and Bradshaw, M.F. (1996). Stereoscopic depth constancy depends on the subject's task. *Vision Research*, **36**, 3441-3456.

Gogel, W.C. (1969). The sensing of retinal size. *Vision Research*, **9**, 3-24.

Gogel, W.C. (1976). An indirect method of measuring perceived distance from familiar size. *Perception & Psychophysics*, **20**, 419-29.

Gogel, W.C. (1977). An indirect measure of perceived distance from oculomotor cues. *Perception and Psychophysics*, **21**, 3-11.

Gogel, W.C. (1978). Size, distance and depth perception. In E.C. Carterette and M.P. Friedman (Eds.), *Handbook of Perception* (Vol. 9, p. 305-308).

Gogel, W.C. (1979). The common occurrence of errors of perceived distance. *Perception and Psychophysics*, **25**, 2-11.

Gogel, W.C. (1980). The sensing of retinal motion. *Perception & Psychophysics*, **28**, 155-163.

Gogel, W.C. (1981). Perceived depth is a necessary factor in apparent motion concomitant with head motion: a reply to Shebliske and Proffitt. *Perception and Psychophysics*, **29**, 173-177.

Gogel, W.C. (1982). Analysis of the perception of motion concomitant with a lateral motion of the head. *Perception & Psychophysics*, **32**, 241-50.

Gogel, W. C. (1983). An illusory motion of a stationary target during head motion is unaffected by paradoxical retinal motion: a reply to Shebilske and Proffitt. *Perception and Psychophysics*, **34**, 482-487.

Gogel, W.C. (1990). A theory of phenomenal geometry and its applications. *Perception & Psychophysics*, **48**, 105-123.

Gogel, W.C. and Tietz, J.D. (1973). Absolute motion parallax and the specific distance tendency. *Perception & Psychophysics*, **13**, 284-292.

Gogel, W.C. and Tietz, J. (1974). The effect of perceived distance on perceived movement. *Perception and Psychophysics*, **16**, 70-78.

Gogel, W.C. and Tietz, J.D. (1977). Eye fixation and attention as modifiers of perceived distance. *Perceptual and Motor Skills*, **45**, 434-462.

Gonzalez, E.G., Steinbach, M.J., Ono, H. and Wolf, M.E. (1989). Depth perception in children enucleated at an early age. *Clinical Vision Sciences*, **4**, 173-177.

Gordon, D.A. (1965). Static and dynamic fields in human space perception. *Journal of the Optical Society of America*, **55**, 1296-1303.

Graham, C.H. (1951). Visual Perception. In S.S. Stevens (Ed.) *Handbook of Experimental Psychology* (p. 868-920). New York: Wiley.

Graham, C.H. (1965). Perception of Movement. In C.H. Graham (ed.), *Vision and Visual Perception*. New York: Wiley.

Graham, C.H., Baker, K.E., Hecht, M. and Lloyd V.V. (1948). Factors influencing thresholds of monocular movement parallax. *Journal of Experimental Psychology*, **38**, 205-223.

Graham, M.E. and Rogers, B.J. (1982a). Simultaneous and successive contrast effects in the perception of depth from motion-parallax and stereoscopic information. *Perception*, **11**, 247-262.

Graham, M.E. and Rogers, B.J. (1982b). Interactions between monocular and binocular depth aftereffects. *IOVS*, **22**, 272.

Green, B.F., Jr. (1961). Figure coherence in the kinetic depth effect. *Journal of Experimental Psychology*, **62**, 272-282.

Green, J. (1889). On certain stereoscopical illusions evoked by prismatic and cylindrical spectacle-glasses. *Transactions of the American Ophthalmology Society*, 449-456. Cited in Howard and Rogers, 1995.

Grosso, E., Sandini, G., Tistarelli, M. (1989). 3-D object reconstruction using stereo and motion. *IEEEE Trans. Sys., Man and Cybern*, **19**, 1465.

Hay, J.C. and Sawyer, S. (1969). Position constancy and binocular convergence. *Perception & Psychophysics*, **5**, 310-312.

Hayashibe, K. (1991). Reversals of visual depth caused by motion parallax. *Perception*, **20**, 17-28.

Hell, W. (1981). Research on monocular depth perception between 1884 and 1914: Influence of the German Unfallversicherungsgesetz and of the jurisdiction of the Reichsversicherungsamt. *Perception*,**10**, 683-694.

Helmholtz, H. von (1866). Cited by Ono, Rivest and Ono (1986).

Helmholtz, Hermann von. (1925). In: J.P.C. (ed.), *Helmholtz's Treatise on Physiological Optics*, vol. 3, 295. Optical Society of America, New York.

Hershberger, W.A., Stewart, M.R. and Laughlin, N.K. (1976). Conflicting motion perspective simulating simultaneous clockwise and counterclockwise rotation in depth *Journal of Experimental Psychology: Human Perception and Performance*, **2**, 174-178.

Hildreth, E.C., Grzywacz, N.M., Adelson, E.H. and Inada, V.K. (1990). The perceptual buildup of three-dimensional structure from motion. *Perception and Psychophysics*, **48**, 19-36.

Hirsch, M.J. and Weymouth, F.W. (1948a). Distance discrimination I. Theoretical consideration. *AMA Archives of Ophthalmology*, **39**, 210-223. Cited in Howard and Rogers (1995).

Hirsch, M.J. and Weymouth, F.W. (1948b). Distance discrimination II. Effect on threshold of lateral separation of the test objects. *AMA Archives of Ophthalmology*, **39**, 224-231. Cited in Howard and Rogers (1995).

Hogervorst, M.A., Bradshaw, M.F. and Eagle, R.A. (1998). Evidence for independent channels in the perception of motion-parallax-defined surfaces. ECVP 1998.

Hogervorst, M. and Eagle, R.A. (1997). Biases in perceived structure from motion explained by a bayesian model. AVA, Surrey.

Hogervorst, M.A. and Eagle, R.A. (1998). Biases in three-dimensional structure-from-motion arise from noise in the early visual system. *Proc. R. Soc. Lond. B.*, **265**, 1587-1593.

Hogervorst, M.A., Kappers, A.M.L, and Koenderink, J.J. (1996). Structure-from-motion: a tolerance analysis. *Perception & Psychophysics*, **58**, 449-459.

Hogervorst, M.A., Kappers, A.M.L. and Koenderink, J.J. (1997). Monocular discrimination of rigidly and nonrigidly moving objects. *Perception & Psychophysics*, **59**, 1266-1279.

Howard, I.P. and Rogers, B.J. (1995). *Binocular Vision and Stereopsis*. Oxford Psychology Series No. 29. Oxford University Press.

Howell, D.C. (1992). *Statistical methods for psychology*. Third edition. Duxbury Press.

Ichikawa, M. and Ono, H. (1996). Depth threshold and motion threshold with different common motion velocity.

Ichikawa, M. and Saida, S. (1996). How is motion disparity integrated with binocular disparity in depth perception? *Perception & Psychophysics*, **58**, 271-282.

Ichikiwa, M. and Ono, H. (1997). Depth perception from motion parallax and different viewing times *IOVS*, 368.

Ittelson, W.H. (1960). *Visual Space Perception*. New York: Springer.

Jansson, G. (1977). Perceived bending and stretching motions from a line of points. *Scandinavian Journal of Psychology*, **18**, 209-215.

Jansson, G. and Börjesson, E. (1969). Perceived direction of rotary motion. *Perception & Psychophysics*, **6**, 19-26.

Jansson, G. and Johansson, G. (1973). Visual perception of bending motion. *Perception*, **2**, 321-326.

Jansson, G. and Runeson, S. (1977). Perceived bending motions from a quadrangle changing form. *Perception*, **6**, 595-600.

Johansson, G. (1973). Visual perception of biological motion and a model for its analyis. *Perception & Psychophysics*, **14**, 201-211.

Johansson, G. (1975). Visual Motion Perception. *Scientific American*, **232**, 76-88.

Johansson, G. (1977). Studies of visual perception of locomotion. *Perception*, **6**, 365-376.

Johns, A.M. (1998). *The Movement of Cyclopean Contours*. Unpublished doctoral thesis, University of Oxford.

Johnston, E.B. (1991). Systematic distortions of shape from stereopsis. *Vision Research*, **31**, 1351-1360.

Johnston, E.B., Cumming, B.G. and Landy, M.S. (1994). Integration of stereopsis and motion shape cues. *Vision Research*, **34**, 2259-2275.

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

Kaiser, M.K. (1990). Angular velocity discrimination. *Perception & Psychophysics*, **47**, 149-156.

Kaplan, G.A. (1969). Kinetic disruption of optical texture: the perception of depth at an edge. *Perception & Psychophysics*, **6**, 193-198.

Kenyon, F.C. (1898). A curious optical illusion connected with an electric fan. *Science*, **8**, 371-372.

Kilpatrick, F.P. and Ittelson, W.H. (1951). Three demonstrations involving the perception of movement. *Journal of Experimental Psychology*, **42**, 394-402.

Koenderink, J.J. (1986). Optic Flow. *Vision Research*, **26**, 161-180.

Koenderink, J.J. (1990). *Solid shape*. MIT Press, Cambridge.

Koenderink, J.J. and van Doorn, A.J. (1975). Invariant properties of the motion parallax field due to the movement of rigid bodies relative to the observer. *Optica Acta*, **22**, 773-791.

Koenderink, J.J. and van Doorn, A.J (1976). Local structure of movement parallax of the plane. *Journal of the Optical Society of America*, **66**, 717-723.

Koenderink, J.J. and van Doorn, A.J. (1981). Exterospecific component of the motion parallax field. *Journal of the Optical Society of America*, **71**, 953-957.

Koenderink, J.J. and van Doorn, A.J. (1986). Depth and shape from differential perspective in the presence of bending deformations. *Journal of the Optical Society of America A*, **3**, 242-249.

Koenderink, J.J. and van Doorn, A.J. (1991). Affine structure from motion. *Journal of the Optical Society of America A*, **75**, 377-385.

Kozlowski, L.T. and Cutting, J.E. (1977). Recogising the sex of a walker from a dynamic point array. *Perception & Psychophysics*, **21**, 575-580.

Lagae, L., Maes, H., Raiguel, S., Xiao, D. and Orban, G.A. (1994). Responses of macaque STS neurons to optic flow components: A comparison of areas MT and MST. *Journal of Neurophysiology*, **71**, 1597-1626.

Lappin, J.S., Doner, J.F. and Kottas, B.L. (1980). Minimal conditions for the visual detection of structure and motion in 3D. *Science*, **209**, 717-719.

Lee, D.N. (1980). The optic flow field: The foundation of vision. *Philosophical Transactions of the Royal Society of London, Series B*, **290**, 169-179.

Lee, D.N. and Young, D.S. (1986). *Gearing action to the environment*. Experimental Brain Research, Series 15. Berlin: Springer-Verlag.

Lind, M. (1996). Perceiving motion and rigid structure from optic flow: a combined weak-perspective and polar-perspective approach. *Perception & Psychophysics*, **58**, 1085-1102.

Lippincott, J.A. (1889). On the binocular metamorphopsia produced by correcting glasses. *AMA Archives of Ophthalmology*, **18**, 18-30. Cited in Howard and Rogers, 1995.

Lipson, M.M., Rogers, B.J. and Ledgeway, T. (1998). Disparity scaling from differential perspective: effects of surface shape and texture. *IOVS*, **39**, 2898.

Liter, J. and Braunstein, M.L. (1994). Judged rotation magnitude in structure-from-motion with and without frontal views. *IOVS*, **35**, p. 1742.

Liter, J.C., Braunstein, M.L. and Hoffmann, D.D. (1992). Perceived relative depth and perceived rotational velocity in structure-from-motion displays. *IOVS*, **33(4)**, 1371.

Liter, J.C., Braunstein, M.L. and Hoffmann, D.D. (1993). Inferring structure from motion in two-view and multi-view displays. *Perception*, **22**, 1441-1465.

Locke, J. *An essay concerning human understanding* (Vol. 1). In A.C. Fraser, Ed., 3rd ed., 1694. Oxford: Clarendon, 1894. (Reprinted, New York: Dover, 1959).

Long, N.R. and Over, R. (1973). Stereoscopic depth aftereffects with random-dot patterns. *Vision Research*, **13**, 1283-1287.

Longuet-Higgins, H.C. (1982). The role of the vertical dimension in stereoscopic vision. *Perception*, **11**, 371-376.

Longuet-Higgins, H.C. and Prazdny, K. (1980). The interpretation of moving retinal images. *Proc. R. Soc., London, Series B*, **208**, 385-387.

Loomis, J.M. and Eby, D.W. (1988). Perceiving structure from motion: Failure of shape constancy. In *Proceedings of the Second International Conference on Compuer Vision* (p. 383 - 91). Washington, DC: Institute of Electrical and Electronic Engineers: Computer Society Press.

Mack, A. (1986). Perceptual aspects of motion in the frontal plane. In R.K. Boff, L. Kaufman and J.P. Thomas (Eds.). *Handbook of perception and human performance* (p. 17-1 - 17-38). New York: Wiley.

Maunsell, J.H.R. and Van Essen, D.C. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. II. Binocular interactions and sensitivity to binocular disparity. *Journal of Neurophysiology*, **49**, 1148-1167.

Mayhew, J.E.W. and Longuet-Higgins, H.C. (1982). A computational model of binocular depth perception. *Nature*, **297**, 376-378.

McConkie, A.B. and Farber, J.M. (1979). Relation between perceived depth and perceived motion in uniform flow fields. *Journal of Experimental Psychology: Human Perception and Performance*, **5**, 501-508.

McKee, S.P. (1981). A local mechanism for differential velocity detection. *Vision Research*, **21**, 491-500.

McKee, S.P. (1983). The spatial requirements for fine stereoacuity. *Vision Research*, **23**, 191-198.

Meese, T.S., Harris M.H. and Freeman, T.C.A. (199). Speed gradients and the perception of surface slant: analysis is 2-dimensional not one dimensional. *Vision Research*, **35**, 2879-2888.

Metzger, W. (1966). Das einaugige tiefensehen. In *Handbuch der Psychologie*, Vol. 1:1, Ch. 14 (Verlag für Psychologie). Cited in Johansson (1973).

Miles, W. (1929). Figure for the windmill illusion. *Journal of General Psychology*, **2**, 143-145.

Miles, W.R. (1931). Movement interpretations of the silhouette of a revolving fan. *American Journal of Psychology*, **43**, 392-405.

Mitchison, G.J. and McKee, S.P. (1990). Mechanisms underlying the anisotropy of stereoscopic tilt perception. *Vision Research*, **30**, 1781-1791.

Musatti, C.L. (1924). Sui fenomeni stereocinetici. *Archivio Italiano di Psicologia* **3**, 105-120. Cited in Braunstein and Andersen, 1984.

Nakayama, K. and Loomis, J.M. (1974). Optical velocity patterns, velocity sensitive neurons and space perception: A hypothesis. *Perception*, **3**, 63-80.

Nakayama, K. and Silverman, G.H. (1988a). The aperture problem - 1. perception of non rigidity and motion detection in translating sinusoidal lines. *Vision Research*, **28**, 739-746.

Nakayama, K. and Silverman, G.H. (1988b). The aperture problem - II. spatial integration of velocity information along contours. *Vision Research*, **28**, 747-753.

Nawrot, M. and Blake, R. (1989). Neural integration of information specifying structure from stereopsis and motion. *Science*, **244**, 716-718.

Nawrot, M. and Blake, R. (1991). The interplay between stereopsis and structure from motion. *Perception & Psychophysics*, **49**, 230-44.

Norman, J.F. and Lappin, J.S. (1992). The detection of surface curvatures defined by optical motion. *Perception & Psychophysics*, **51**, 386-396.

Norman, J.F., Lappin, J.S. and Zucker, W.Z. (1991). The discriminability of smooth stereoscopic surfaces.*Perception*, **20**, 789-807.

Norman, J.F. and Todd, J.T. (1993). The perceptual analysis of structure from motion for rotating objects undergoing affine stretching transformations. *Perception & Psychophysics*, **53**, 279-291.

Norman, J. F. and Todd, J.T. (1994). Perception of rigid motion in depth from the optical deformations of shadows and occlusion boundaries. *Journal of Experimental Psychology: Human Perception and Performance*, **20**, 343-356.

Norman, J.F. and Todd, J.T. (1995). The perception of 3-D structure from contradictory optical patterns. *Perception*, **57**, 826-834.

O'Leary, A. amd Wallach, H. (1980). Familiar size and linear perspective as distance cues in stereoscopic depth constancy. *Perception & Psychophysics*, **27**, 131-135.

Ogle, K.N. (1938). Induced size effect I. A new phenomenon in binocular space perception associated with the relative sizes of the images of the two eyes. *AMA Archives of Ophthalmology*, **20**, 604-623. Cited in Howard and Rogers, 1995.

Ohtsuka, S., Ujike, H., Sato, T. and Saida, S. (1997). Depth and motion thresholds from relative image motion. *IOVS*, 38, 393.

Ono, M.E., Rivest, J. and Ono, H. (1986). Depth perception as a function of motion parallax and absolute distance information. *Journal of Experimental Psychology: Human Perception and Performance*, **12**, 331-337.

Ono, H., Rogers, B.J., Ohmi, M. and Ono, M. (1988). Dynamic occlusion and motion parallax in depth perception. *Perception*, **17**, 255-66.

Ono, H. and Steinbach, M.J. (1990). Monocular stereopsis with and without head movement. *Perception & Psychophysics*, **48**, 179-187.

Ono, H. and Ujike, H. (1993). Equal-depth contours as a function of direct velocities of head-movement. ECVP, 81.

Ono, H. and Ujike, H. (1993). Zone in which motion parallax is completely effective. *IOVS*, **34**, 1052.

Orban, G.A., Lagae, L., Verri, A., Raiguel, S., Xiao, D., Maes, H. and Torre, V. (1992). First-order analysis of optical flow in monkey brain. *Proceedings of the National Academy of Science USA*, **89**, 2595-2599.

Owens, D.A. and Leibowitz, H.W. (1980). Accommodation convergence and distance perception in low illumination. *American Journal of Physiological Optics*, **57**, 540-550.

Park, J.N (1964). *The role of motion parallax in depth perception*. (Doctoral dissertation, University of Kansas). Ann Arbor, Mich.: University Microfilms, 1964, no. 65-7668.

Patterson, R.D. (1976). Auditory filter shapes derived with noise stimuli. *Journal of the Acoustical Society of America*, **58**, 640-654.

Perotti, V.J., Todd, J.T. and Farley Norman, J. (1996). The visual perception of rigid motion from constant flow fields. *Perception & Psychophysics*, **58**, 666-679.

Perrone, J.A. (1992). Model for the computation of self-motion in biological systems. *Journal of the Optical Society of America A*, **9**, 177-194.

Petersik, J.T. (1980a). Rotation judgements and depth judgements: separate or dependent processes? *Perception & Psychophysics*, **27**, 588-590.

Petersik, J.T. (1980b). The effects of spatial and temporal factors on the perception of stroboscopic rotation simulations. *Perception*, **9**, 271-283.

Petersik, J.T. (1991). Perception of three-dimensional angular rotation. *Perception & Psychophysics*, **50**, 465-474.

Petersik, J.T., Shepard, A. and Malsch, R. (1984). A three-dimensional motion aftereffect produced by prolonged adaptation to a rotation simulation. *Perception*, **13**, 489-497.

Poggio, G.F. and Talbot, W.H. (1981). Mechansims of static and dynamic stereopsis in foveal cortex of the rhesus monkey. *Journal of Physiology*, **315**, 469-492.

Pollick, F.E. (1994). Perceiving shape from profiles. *Perception & Psychophysics*, **55**, 152-161.

Pollick, F.E. (1997). The perception of motion and structure in structure-from-motion: comparisons of affine and euclidean formulations. *Vision Research*, **37**, 447-466.

Pollick, F.E., Nishida, S., Koike, Y. and Kawato, M. (1994). Perceived motion in structure from motion: pointing responses to the axis of rotation. *Perception & Psychophysics*, **56**, 91-109.

Pollick, F.E. and Wilson, L.L. (1991).The perception of shape and angular velocity from the profiles of a rotating solid. *IOVS*, **32**, 1277.

Power, R.P. (1967). Stimulus properties which reduce apparent reversal of rotating rectangular shapes.*Journal of Experimental Psychology*, **73**, 595-599.

Power, R.P. and Day, R.H. (1973). Constancy and illusion of apparent direction of rotary motion in depth: tests of a theory. *Perception & Psychophysics*, **13**, 217-223.

Predebon, J. (1993). The familiar-size cue to distance and stereoscopic depth perception. *Perception*, **22**, 985-995.

Prince, S.J.D. (1998). *Measurement and encoding of binocular disparity in human stereopsis*. Unpublished doctoral thesis, University of Oxford.

Prince, S.J.D. and Rogers, B.J. (1998). Sensitivity to disparity corrugations in peripheral vision. *Vision Research*, **38**, 2533-2537.

Proffitt, D.R., Rock, I., Hecht, H. and Schubert, J. (1992). Stereokinetic effect and its relation to the kinetic depth effect. *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 3-21.

Ramachandran, V.S., Cobb, S. and Rogers-Ramachandran, D. (1988). Perception of 3-D structure from motion: the role of velocity radients and segmentation boundaries. *Perception & Psychophysics*, **44**, 390-393.

Regan, D. and Beverley, K.I. (1978). Looming detectors in the human visual pathway. *Vision Research*, **18**, 415-421.

Reppas, J.B., Niyogi, S., Dale, A.M., Sereno, M.I. and Tootell, R.B.H. (1997). *Nature*, **388**, 175-179.

Richards, W. (1985). Structure from stereo and motion. *Journal of the Optical Society of America A*, 343-349.

Rivest, J., Ono, H. and Saida, S. (1989). The roles of convergence and apparent distance in depth constancy with motion parallax. *Perception and Psychophysics*, **46**, 401-408.

Rock, I. (1984). *Perception*. New York, Scientific American Library.

Rock, I. (1995). *Perception* (Rev. Ed.). New York, Scientific American Library.

Rogers, B.J. (1988). Perspectives on movement. *Nature*, **333**, 16-17.

Rogers, B.J. (1993). Motion parallax and other dynamic cues for depth in humans. In *Visual motion and its role in the stabilization of gaze*. (ed. F.A. Miles and J. Wallman), p. 119-137. Elsevier, Amsterdam.

Rogers, B.J. and Bradshaw, M.F. (1991). Vertical disparities in binocular and monocular depth perception. *Perception*, **20**, 130A.

Rogers, B.J. and Bradshaw, M.F. (1992). Differential perspective effects in binocular stereopsis and monocular depth perception. *IOVS*, **33**, 1333.

Rogers, B.J. and Bradshaw, M.F. (1993). Vertical disparities, differential perspective and binocular stereopsis. *Nature*, **361**, 253-255.

Rogers, B.J. and Collett, T.S. (1989). The appearance of surfaces specified by motion parallax and binocular disparity. *Quarterly Journal of Experimental Psychology*, **41A**, 697-717.

Rogers, B.J. and Graham, M. (1979). Motion parallax as an independent cue for depth perception. *Perception*, **8**, 125-134.

Rogers, B.J. and Graham, M. (1982). Similarities between motion parallax and stereopsis in human depth perception. *Vision Reaearch*, **22**, 261-70.

Rogers, B.J. and Graham, M.E. (1983). Anisotropies in the perception of three-dimensional surfaces. *Science*, **221**, 1409-1411.

Rogers, B.J. and Graham, M.E. (1985). Motion parallax and the perception of three-dimensional surfaces. In *Brain mechanisms and spatial vision*. Eds. D. Ingle, M. Jeannerod and D. Lee, p. 95-111. Martinus Nijhoff, The Hague.

Rogers, B.J. and Koenderink, J. (1986). Monocular aniseikonia: A motion parallax analogue of the disparity-induced effect. *Nature*, **322**, 62-63.

Rogers, S. and Rogers, B.J. (1992). Visual and nonvisual information disambiguate surfaces specified by motion parallax. *Perception and Psychophysics*, **52**, 446-452.

Rouse, M.W., Tittle, J.S. and Braunstein, M.L. (1989). Stereoscopic depth perception by static stereo-deficient observers in dynamic displays with constant and changing disparity. *Optometry and Vision Science*, **65**, 355-362.

Rovamo, J. and Virsu, V. (1979). An estimation and application of the human cortical magnification factor. *Experimental Brain Research*, **37**, 495-510.

Roy, J.P., Komatsu, H. and Wurtz, R.H. 1992). Disparity sensitivity of neurons in monkey extrastriate area MST. *Journal of Neuroscience*, **12**, 2478-2492.

Ryan, C. and Gillam, B. (1994). Cue conflict and stereoscopic surface slant about horizontal and vertical axes. *Perception*, **23**, 645-658.

Sachtler, W.L. and Zaidi, Q. (1995). Visual processing of motion boundaries. *Vision Research*, **35**, 807-826.

Saida, S. and Ono, H. (1989). What extent of parallax produces depth perception? *IOVS*, **30**, 87.

Saidpour, A., Braunstein, M.L. and Hoffman, D.D. (1993). Interpolation across surface discontinuities in structure from motion. *Perception & Psychophysics*, **55**, 611-622.

Sakurai, K. and Ono, H. (1989). Perceived depth from radial optic flow yoked to forward and backward head movement. *IOVS*, **30**, 253.

Sakuri, K. and Ono, H. (1996). Perceived depth from dot density change by lateral head movements. ECVP Abstracts.

Sakuri, K. and Ono, H. (1997). Perceived depth from expansion/contraction and shearing motion yoked to lateral and vertical head movements. *IOVS*, **38**, 395.

Sato, T. (1989). Personal communication to Ono and Steinbach. Cited in Ono and Steinbach (1990).

Schiff, W. and Detwiler, M.L. (1979). Information used in judging impending collision. *Perception*, **8**, 647-658.

Schiff, W. Caviness, J.A. and Gibson, J.J. (1962). Persistent fear responses in rhesus mokeys to the optical stimulus of "looming". *Science*, **136**, 982-983.

Schumer, R. (1979). *Mechanisms in Human Sterepsis*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.

Schumer, R.A. and Ganz, L. (1979). Independent stereoscopic channels for different extents of spatial pooling. *Vision Research*, **19**, 1303-1314.

Shebilske, W.L. and Proffitt, D.R. (1981). The priority of perceived distance for perceiving motion has not been demonstrated: critical comments on Gogel's "The sensing of retinal motion"*Perception and Psychophysics*, **29**, 170-172.

Shebilske, W.L. and Proffitt, D.R. (1983). Paradoxical retinal motions during head movements: Apparent motion without equivalent apparent displacement. *Perception and Psychophysics*, **34**, 476-481.

Smith, O.W. and Smith, P.C. (1963). On motion parallax and perceived depth. *Journal of Experimental Psychology*, **65**, 107-108.

Snippe, H.P. and Werkhoven, P. (1993). Pulse modulation detection in human motion vision. *Vision Research*, **33**, 907-914.

Snowden, R.J. and Braddick, O.J. (1991). The temporal integration and resolution of velocity signals. *Vision Research*, **31**, 907-914.

Sobel, E.C. and Collett, T.S. (1991). Does vertical dispariy scale the perception of stereoscopic depth? *Proceedings of the Royal Society, London*, **B244**, 87-90.

Sparrow, J.E. and Stine, W.M.W. (1998). The perceived rigidity of rotating eight-vertex geometric forms: extracting nonrigid structure from rigid motion. *Vision Research*, **38**, 541-556.

Steinbach, M.J., Ono, H. and Wolf, M.E. (1991). Motion parallax judgements of depth as a function of the direction and type of head movement. *Canadian Journal of Psychology*, **45**, 92-98.

Stevens, K.A. and Brookes, A. (1988). Integrating stereopsis with monocular interpretations of planar surfaces. *Vision Research*, **17**, 35-42.

Swanston, M.T. and Gogel, W.C. (1986). Perceived size and motion in depth from optical expansion. *Perception & Psychophysics*, **39**, 309-326.

Swanston, M.T. and Wade, N.J. (1988). The perception of visual motion during movements of the eyes and of the head. *Perception & Psychophysics*, **43**, 559-566.

Swanston, M.T., Wade, N.J., Ono, H. and Shibuta, K. (1992). The interaction of perceived distance with the perceived direction of visual motion during movements of the eyes and the head. *Perception & Psychophysics*, **52**, 705-713.

Swash, S.A. and Rogers, B.J. (1996). The perceived rotation of motion parallax surfaces: dependence on spatial frequency of depth corrugations. *IOVS*, **37**, 1317.

Swash, S.A. and Rogers, B.J. (1998). Transparency in motion parallax and stereopsis: similarities and differences. *IOVS*, **39**, 3950.

Tanaka, K., Fukada, Y. and Saito, H. (1989). Underlying mechanisms of the response specificity of expansion/contraction and rotation cells in the dorsal part of the medial superior temporal area of the macaque monkey. *Journal of Neurophysiology*, **62**, 642-656.

Tittle, J.S. and Braunstein, M.L. (1991). Shape perception from binocular disparity and structure-from-motion. In P.S. Schenker (ed.). *Sensor Fusion iii: 3-D Perception and Recognition. Proceedings of the SPIE.* **1383**, 225-234.

Tittle, J.S. and Braunstein, M.L. (1993). Recovery of 3-D shape from binocular disparity and structure from motion. *Perception & Psychophysics*, **54**, 157-169.

Tittle, J.S. and Perotti, V.J. (1997). The perception of shape and curvedness from binocular stereopsis and structure from motion. *Perception & Psychophysics*, **59**, 1167-1179.

Tittle, J.S., Todd, J.T., Perotti, V.J. and Norman, J.F. (1995). Systematic distortion of perceived three-dimensional structure from motion and binocular stereopsis. *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 663-678.

Todd, J.T. (1982). Visual information about rigid and nonrigid motion: a geometric analysis. *Journal of Experimental Psychology: Human Perception and Performance*, **8**, 238-252.

Todd, J.T. (1983). The perception of gait. *Journal of Experimental Psychology: Human Perception and Performance*, **9**, 31-42.

Todd, J.T. (1984). The perception of three-dimensional structure from rigid and non-rigid motion. *Perception & Psychophysics*, **36**, 97-103.

Todd, J.T. and Akerstrom, R.A. (1987). Perception of three-dimensional form from patterns of optical texture. *Journal of Experimental Psychology: Human Perception and Performance*, **13**, 242-255.

Todd, J.T. and Bressan, P. (1990). The perception of 3-dimensional affine structure from minimal apparent motion sequences. *Perception &Psychophysics*, **48**, 419-430.

Todd, J.T. and Norman, J.F. (1991). The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Perception & Psychophysics*, **50**, 509-523.

Tschermak-Seysenegg, A. (1939). Über Parallaktoskopie. *Pflugers Archiv Physiologie* **241**, 454-469. Cited in Rogers (1993).

Turner, J. and Braunstein, M.L (1995). Size constancy in structure from motion. *Perception*, **24**, 1155-1164.

Turner, J., Andersen, G.J. and Braunstein, M.L. (1995). Detection of surfaces in structure from motion. *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 809-821.

Turner, J., Braunstein, M.L. and Andersen, G.J. (1997). Relationship between binocular disparity and motion parallax in surface detection. *Perception & Psychophysics*, **59**, 370-380.

Tyler, C.W. (1975). Spatial organisation of binocular disparity sensitivity. *Vision Research*, **15**, 583-590.

Tyler, C.W. (1983). Sensory processing of binocular disparity. In *Vergence eye movements: Basic and clinical aspects*, (ed. M. C Schor and K. J. Ciuffreda), p. 199-296. Butterworth, Boston.

Ujike, H. and Ono, H. (1992). Parallactic depth threshold with different amplitudes and frequencies of head movement. *IOVS*, **33**, 1374.

Ujike, H. and Ono, H. (1993). Displacement signal is used for parallactic depth. *IOVS*, 1190.

Ullman, S. (1979). *The interpretation of visual motion*. MIT Press, Cambridge, MA.

Ullman, S. (1983). The measurement fo visual motion: computational considerations and some neuropsychological implications. *Trends in Neurosciences*, **6**, 177-179.

van Damme, W.J.M. and van de Grind, W.A. (1993). Active vision and the identification of three-dimensional shape. *Vision Research*, **33**, 1581-1587.

van Damme, W.J.M. and van de Grind, W.A. (1996). Non visual information from SFM. *Vision Research*, **36**, 3119-3127.

van Damme, W.J.M., Oosterhoff, F.H. and van de Grind, W.A. (1994). Discrimination of 3-D shape and 3-D curvature from motion in active vision. *Perception & Psychophysics*, **55**, 340-349.

van Veen, H.A.C. and Werkhoven, P. (1996). Metamerisms in structure-from-motion perception. *Vision Research*, **36**, 2197-2210.

von Fieandt, K. and Gibson, J.J. (1959). The sensitivity of the eye to two kinds of continuous transformation of a shadow pattern. *Journal of Experimental Psychology*, **57**, 344-347.

Wallach, H. (1985a). Perceiving a stable environment. *Scientific American*, **252**, 92-98.

Wallach, H. (1985b). Learned stimulation in space and motion perception. *American Psychologist*, **40**, 399-404.

Wallach, H. (1987). Perceiving a stable environment when one moves. *Annual Review of Psychology*, **38**, 1-27.

Wallach, H. and Flaherty, E.W. (1975). A compensation for field expansion caused by moving forward. *Perception & Psychophysics*, **17**, 445-449.

Wallach, H. and O'Connell, D.N. (1953). The kinetic depth effect. *Journal of Experimental Psychology*, **45**, 205-17.

Wallach, H. and Zuckerman, C. (1963). The constancy of stereoscopic depth. *American Journal of Psychology*, **76**, 404-412.

Wallach, H., Bacon, J. and Schulman, P. (1978). Adaptation in motion perception: Alteration of induced motion. *Perception & Psychophysics*, **24**, 509-514.

Wallach, H., Weisz, A. and Adams, P.A. (1956). Circles and derived figures in rotation. *American Journal of Psychology*, **69**, 48-59.

Wallach, H., Yablick, G.S. and Smith, A. (1972). Target distance and adaptation in distance perception in the constancy of visual direction. *Perception & Psychophysics*, **12**, 139-145.

Welch, L. and Tolstoshev, N.G. (1996). Is shearing motion detected by spatial frequency tuned mechanisms? *IOVS*, **37**, 4218; Personal communication.

Welch, L. (1996). Pers. comm.

Werkhoven, P. and van Veen, H.A.H.C. (1995). Extraction of relief from visual motion. *Perception & Psychophysics*, **57**, 645-656.

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

Whipple, W.R. and Wallach, H. (1978). Direction-specific motion thresholds for abnormal image shifts during saccadic eye movement. *Perception & Psychophysics*, **24**, 349-355.

Williams, J.S. (1993). *An experimental investigation of depth cue interaction*. Unpublished DPhil thesis, Oxford University.

Youngs, W.M. (1976). The influences of perspective and disparity on the perception of slant. *Vision Research*, **16**, 79-82.

Yuille, A., Bülthoff, H. and Fahle, M. (1987). Theoretical model of disparity gradients. *Journal of the Optical Society of America A*, **4**, 14.

Zegers, R.T. (1948). Monocular movement parallax thresholds as functions of field size, field position and speed of stimulus movement. *Journal of Physiology*, **26**, 477-498.

Zegers, R.T. (1964). The reversal illusion of the Ames trapezoid. *Transactions of the New York Academy of Sciences*, **26**, 377-400.

Zihl, J., Von Cramon, D. and Mai, N. (1983). Selective disturbance of movement vision after bilateral brain damage. *Brain*, **166**, 313-340.