A Perturbation Analysis of Depth Perception from Combinations of Texture and Motion Cues

MARK J. YOUNG,*† MICHAEL S. LANDY,* LAURENCE T. MALONEY*

Received 23 April 1992; in revised form 3 May 1993

We examined how depth information from two different cue types (object motion and texture gradient) is integrated into a single estimate in human vision. Two critical assumptions of a recent model of depth cue combination (termed *modified weak fusion*) were tested. The first assumption is that the overall depth estimate is a weighted linear combination of the estimates derived from the individual cues, after initial processing needed to bring them to a common format. The second assumption is that the weight assigned to a cue reflects the apparent reliability of that cue in a particular scene. By this account, the depth combination rule is linear and dynamic, changing in a predictable fashion in response to the particular scene and viewing conditions. A novel procedure was used to measure the weights assigned to the texture and motion cues across experimental conditions. This procedure uses a type of *perturbation analysis*. The results are consistent with the weighted linear combination rule. In addition, when either cue is corrupted by added noise, the weighted linear combination rule shifts in favor of the uncontaminated cue.

Depth Multiple cues Sensor fusion

INTRODUCTION

There are a variety of measures in two-dimensional (2D) images which provide information about threedimensional (3D) structure. Psychophysical studies of biological depth perception have most commonly studied these depth cues in isolation: depth from binocular disparity (Julesz, 1971; Mayhew & Frisby, 1981), depth from texture gradient (Gibson, 1950; Todd & Akerstrom, 1987), structure from motion (Wallach & O'Connell, 1953; Sperling, Landy, Dosher & Perkins, 1989), and others. These studies leave open the problem of how depth information from different cue types is integrated into a single estimate at each location in a scene. This problem is receiving an increasing amount of attention (Bülthoff & Mallot, 1988; Aloimonos & Shulman, 1989; Maloney & Landy, 1989).

Clark and Yuille (1990) divide approaches to the problem into *weak fusion* and *strong fusion* models. Research on single depth cues has fostered an understanding of depth perception as a system of modular processes (Marr, 1982), and weak fusion models emphasize the apparently modular structure of depth vision. They assume that even when multiple depth cues are simultaneously present in a scene, each type of depth cue is processed independently to produce an estimate of depth. These separate depth estimates are then combined at each location by a rule of combination.

In contrast, strong fusion models assume that information from multiple cue types is processed cooperatively to arrive at a single depth estimate. By this account, the apparently modular structure of the visual system is an artifact of the kinds of experiments used to investigate it. Single cue experiments may only demonstrate that depth perception is organized to accommodate missing information. Strong fusion models formulate depth estimation as a nonseparable function of multiple cues and predict nonlinear interactions between depth information derived from different types of cues. Weak fusion models preclude such interactions.

Studies of cues strongly in conflict have demonstrated complex interactions between different types of depth information (Wallach & Karsh, 1963; Epstein, 1968; Johnston, Cumming & Parker, 1993) and, consequently, the simple weak fusion models described above cannot be valid descriptions of human visual processing. The rejection of these models leads naturally to a search for a model of depth processing that parsimoniously accounts for human depth estimation in scenes where multiple types of cues are present, and which also explains the interactions observed between cue types.

Maloney and Landy (1989) (also see Landy, Maloney, Johnston & Young, 1991b) have proposed a framework for investigating depth cue combination

^{*}Psychology Department and Center for Neural Science, New York University, 6 Washington Place, Room 868, New York, NY 10003, U.S.A.

[†]To whom all correspondence should be addressed.



FIGURE 1. Illustration of the manner in which depth from texture (d_t) and depth from motion (d_m) are varied independently. Round texture elements of random size and spacing are formed by the intersections of a cylinder (with depth d_t) with randomly-placed balls. The resulting textured surface is then projected (parallel to z-axis) onto a second cylinder (with depth d_m). Stimuli are generated by projections (parallel to z-axis) of the second surface rotated about the x-axis. The width of the projected surface is 2 W.

which they term *modified weak fusion*. They suggest that in natural scenes, individual depth cues tend towards consonance and may be combined in a normative fashion. They assume that depth estimates derived from single cue types are combined by a weighted average, consistent with a weak fusion approach. This choice of rule of combination is consistent with some previous research. For example, Dosher, Sperling and Wurst (1986) found that a weighted linear combination explained their data for combinations of binocular disparity and proximity-luminance covariance ("brighter is closer") cues in conjunction with the kinetic depth effect (KDE). Maloney and Landy assume that interactions between cues are due to two normative processes they term (1) cue promotion and (2) dynamic reweighting.

Cue promotion

Single cues may provide depth information dependent on one or more undetermined parameters (e.g. Richards, 1985). Depth from object motion (KDE), for example, does not provide estimates of absolute depth to points of an object, but instead provides estimates which are missing two unknown parameters: the distance to any single point on the object and the direction of object rotation. Maloney and Landy (1989) argue that "depth-maps-with-parameters" are first computed in a modular fashion for each cue type. The resulting estimates are then *promoted*: the missing parameters in each depth map are filled in by comparison with others. The process of cue promotion constitutes an interaction among different cues.

Dynamic reweighting

The promoted depth cues are then combined by a weighted average. If the weights are chosen correctly, the resulting depth estimate will be no worse than, and likely better than, any of the individual depth estimates. Maloney and Landy argue that the weight assigned to a depth estimate derived from a particular cue type should reflect the estimated reliability of that cue type in the current scene, under the current viewing conditions. The weight assigned to depth estimates derived from motion parallax, for example, should depend on the observer's velocity as estimated from both visual and nonvisual (e.g. vestibular) information. A normative system should use such *ancillary measures* to estimate the reliability of cues, and to determine appropriate weights.

This paper reports work testing the assumptions of the framework proposed by Maloney and Landy. We have tested the linearity of observers' depth estimates for combinations of texture and motion cues by a psychophysical method inspired by perturbation analysis. We analyzed changes in depth perceived from two consonant cues as a function of a small difference between the cues. This work is a continuation of preliminary work reported by Landy, Maloney and Young (1991a). In a stimulus containing both texture and motion cues to depth, the depth of a stimulus as portrayed by texture (d_1) and the depth portrayed by motion (d_m) are manipulated independently. To measure the combined depth estimate for a stimulus with inconsistent texture and motion cues $(d_1 \neq d_m)$, it is determined which stimulus with consistent cues $(d'_{t} = d'_{m} = d)$ appears to have equal depth, thus empirically specifying the function which yields the combined depth estimate for independent texture and motion cues,

$$d = f(d_{\rm t}, d_{\rm m}). \tag{1}$$

If the combined depth estimate is a weighted average of the individual cues,* then

$$d = \alpha_{\rm t} d_{\rm t} + \alpha_{\rm m} d_{\rm m}, \qquad (2)$$

where the weight of the texture cue (α_t) and the weight of the motion cue (α_m) sum to 1

$$\alpha_{\rm t} + \alpha_{\rm m} = 1. \tag{3}$$

If the combined depth estimate is linear in d_t , the weight assigned to the texture cue can be measured by measuring the combined depth estimate d while we vary the depth portrayed by texture d_t , holding the depth portrayed by motion d_m fixed. For two measurements,

$$d_1 = \alpha_t d_{t_1} + \alpha_m d_m, \qquad (4)$$

^{*}By using the depth cues of object motion and surface texture, we skirt around the issue of cue promotion. Both motion and texture can provide "depth-maps-with-parameters", but they share the same parameterization. The depth provided by each cue scales with the viewing distance, and can additionally involve a reflection about the image plane (a depth reversal). Here, we assume only that texture and motion share the same estimates of viewing distance and depth order.



FIGURE 2. Example stimuli, labeled by eccentricity (d_i/W) . The range shown is larger than that used in the experiments.

and

$$d_2 = \alpha_t d_{t_2} + \alpha_m d_m, \qquad (5)$$

we then have

$$\alpha_{t} = \frac{d_2 - d_1}{d_{t_2} - d_{t_1}} = \frac{\Delta d}{\Delta d_t}.$$
 (6)

Note that this derivation does not depend on the assumption that the weights sum to 1 [equation (3)], but simply that α_m remains constant over the two levels of d_t . The weight of the motion cue can be measured in a similar way. In this study we tested the linear combination rule by determining if the combined depth estimate was a linear function of each cue (Expt 1). The role of ancillary information in determining the weights was examined by testing cue combinations at different levels of single cue reliability (Expt 2).

EXPERIMENT 1

Methods

Stimuli. The stimuli were motion sequences portraying a pair of rotating surfaces. Each surface was an elliptic hemicylinder covered with approximately circular texture elements (Fig. 1). The axis of rotation was horizontally oriented in a frontoparallel plane. A cylinder's elliptic cross section had one semiaxis that was always parallel to the image plane, corresponding to the cylinder's halfwidth (W), which was constant across all stimuli. The other semiaxis was parallel to the line of sight when the cylinder's center axis was in a frontoparallel plane (i.e. vertical). This semiaxis corresponded to the cylinder's depth, which was varied across stimuli. Two depth values, d_t and d_m , were used to construct each stimulus surface.

A homogeneous and isotropic texture was mapped to a surface comprised of points with coordinates (x, y, z)where $(x/W)^2 + (z/d_1)^2 = 1$. The depth coordinates (z) of this surface were then rescaled $z \rightarrow (d_m/d_t)z$ so that surface points then had coordinates where $(x/W)^2 + (z/d_m)^2 = 1$. A stimulus view was rendered from a parallel projection of the rescaled cylinder rotated about the x-axis with the line of sight along the z-axis. The projected texture compression was that of a homogeneously and isotropically textured cylinder with eccentricity d_t/W . The projected motion was that of an inhomogeneously and anisotropically textured cylinder with eccentricity d_m/W . Shape-from-texture algorithms generally begin with the assumption that the gradient of texture compression is due to the changing surface orientation and not changes in the physical texture distribution. Adopting this assumption, the stimulus generation technique just described allows us to independently vary d_t and d_m (different values of d_t are illustrated in Fig. 2). The perturbation analysis should be limited to differences between d_t and d_m that do not disturb this assumption by providing contrary information to the observer.

The 3D texture function was defined procedurally,* based upon a volume randomly packed to density 0.25 with nonintersecting spheres 20 pixels in diameter. Texture elements were shaded black on a white background. All cylinders were 240 pixels wide (2W) with eccentricities $(d_t/W, d_m/W)$ of 0.83, 1.04, 1.25, 1.46, or 1.67. Motion sequences were projections of a cylinder rotated about the x-axis from -15 to +15 deg (with respect to vertical) in 0.625 deg steps (49 frames/surface). The viewing window was 240 horizontal \times 256 vertical pixels

^{*}A texture volume is constructed by randomly packing a virtual box with spheres. The box is positioned in the coordinate frame of a surface so as to contain the surface. The color of a surface pixel is that of the voxel (primarily) containing it. The color of a voxel depends upon its position within the texture volume: voxels (primarily) within spheres are black and voxels not within spheres are white. The surface is thus marked with randomly positioned nearly circular texture elements of random size. With respect to surface area, both the size and density of texture elements are homogeneously and isotropically distributed. The general form of this texture model is termed *bombing* by Peachey (1985). The "box of balls" implementation used in these experiments is related to the method used by Johnston *et al.* (1993).

with the rotation axis at midheight in the window. The vertical length of a cylinder was such that the viewing window always clipped the bottom and top portions, eliminating the depth cue provided by the time-varying deformations of the contours at the edges of the surface.

To insure that subjects were making judgments based upon depth estimates rather than 2D information such as single-point displacement magnitudes and velocities, we reduced the correlation between d_m and such 2D measures. For each surface on each trial, the rate and total amplitude of rotation were randomized. The rotation rate was randomly sampled from the range 12.5-37.5 deg/sec with uniform probability.* The rotation amplitude was randomly sampled from the range 10-30 deg (symmetric about the line of sight) with uniform probability. A stimulus display began with the downward-most rotated frame and then rotated upward, then back down, and so on. This cycle was repeated as many times as full rotations could be shown in 4.85 sec, with remaining frames blank (black). Three full rotations were displayed on average.

A complete stimulus sequence was constructed by displaying two such surfaces side by side on a black background with rotation rates and amplitudes independently sampled. Each displayed surface window subtended 3.5 deg of visual angle and was viewed from a distance of 1.7 m. The two windows were separated by 15 min of visual angle.

Stimuli were displayed using an Adage RDS-3000 display system driving a US Pixel PX15 monochrome monitor with a fast P4-like phosphor. The video format was 60 Hz, noninterlace, with 480×512 visible pixels. The background luminance was 32 cd/m^2 when a stimulus was displayed. Displays were viewed monocularly through a reduction screen approximately matched in color and luminance to the background of the stimuli.

The perceived dimensions of the cylinders depended on observer calibration of viewing distance (e.g. Ono, Rivest & Ono, 1986) which we tried to eliminate with the reduction screen.[†] For references to stimulus depths however, instead of pixel numbers we use a metric scale consistent with the viewing distance employed in the experiments (0.044 cm/pixel). The depths reported would be consistent with cylinders 10.5 cm wide and texture volumes packed with spheres with 8.8 mm diameters.

Procedure. In each trial, two surfaces were displayed side by side. One of the surfaces was a consistent-cues surface where $d_t = d_m$. The other surface was a mixed-cues surface where generally $d_t \neq d_m$. The subject's task was to indicate with a key press which surface appeared to extend further in depth. The side on which a surface appeared (right or left) was randomly determined on each trial. A response terminated the display and the next trial would begin after a 0.5 sec pause. Subjects were given the option of not responding, in which case the stimulus display would be presented again after a 0.5 sec delay.

Within a block of trials there was a single mixed-cues surface. A block consisted of 50 two-alternative forcedchoice (2AFC) comparisons with each of five consistentcues surfaces in random order for 250 trials per block. Consistent-cues surfaces portrayed depths of 4.4, 5.5, 6.6, 7.7, and 8.8 cm. For half of the blocks, the mixedcues surface had d_t fixed at 6.6 cm while d_m was 4.4, 5.5, 6.6, 7.7, or 8.8 cm. On the other half of the blocks, the mixed-cues surface had d_m fixed at 6.6 cm while d_t was 4.4, 5.5, 6.6, 7.7, or 8.8. cm. There were two blocks for



Consistent Cues Depth (cm)

FIGURE 3. The psychometric functions of the proportion of comparisons for which the consistent-cues surface was perceived to have greater depth than the mixed-cues surface. The depth portrayed by the consistent-cues surface varies along the abscissa. Each curve represents the comparisons with a single mixed-cues surface and combines the data from two blocks of 250 trials. Each data point is based on 100 comparisons. Each row corresponds to a different subject. The lefthand column plots data for blocks in which the depth portrayed by motion (d_m) in the mixed-cues stimulus was fixed and depth portrayed by texture (d_i) varied from block to block; in the right-hand column d_t was fixed and d_m varied across blocks. In both cases, the curves appear to slide rightward as the depth portrayed by the variable cue is increased.

^{*}Equipment limitations required an overall sampling rate (new frames/sec) of 20 Hz. Thus, every 0.05 sec the current orientation of the surface was determined. The closest available frame, given the 0.625 deg sampling density of rotation, was then displayed for three full video frames (at a 60 Hz video frame rate). At the slowest rotation rate of 12.5 deg/sec, all 49 frames of the precomputed sequence were simply shown in order at 20 Hz. At the fastest rate, they were subsampled at a rate of 3:1. For intermediate rates, the precomputed frames were subsampled as necessary to best approximate the targeted rate, which often resulted in nonuniform sampling of rotations. This resulted in a minor oscillation in the actual rate about the targeted rate. However, the rotational motion was always smooth in appearance.

[†]The primary function of the reduction screen was to provide a monocular view of a blank field (except for the surface stimuli). To view through the reduction screen, an observer was required to position their head against a mask that restrained lateral head movements. This served to eliminate motion parallax cues.

each mixed-cues surface for 20 blocks total and 100 trials per data point. The order of blocks was randomized.

Subjects. Subjects MJY and LTM were authors. Subject EBJ was knowledgeable about the purposes of the experiment. All subjects had normal or corrected-tonormal vision. All subjects were experienced psychophysical observers.

Results and discussion

Figure 3 graphs the psychometric functions of the proportion of comparisons for which the consistent-cues surface was perceived to have greater depth than the mixed-cues surface. The depth portrayed by the consistent-cues surface varies along the abscissa. Each curve represents the comparisons with a single mixed-cues surface and combines the data from two blocks of 250 trials. Each data point is based on 100 comparisons. Each row of Fig. 3 graphs the data for a single subject over 20 blocks of trials. The psychometric functions in the left columns of Fig. 3 represent the comparisons with mixed-cues surfaces of fixed d_m (6.6 cm) with d_t varied (4.4, 5.5, 6.6, 7.7, and 8.8. cm). The psychometric functions in the right columns of Fig. 3 represent the comparisons with mixed-cues surfaces of fixed d_1 (6.6 cm) with $d_{\rm m}$ varied (4.4, 5.5, 6.6, 7.7, and 8.8 cm). In the following analysis, each panel represents one of the six sets of subject \times condition.

In each plot the central psychometric function is for a so-called mixed-cues surface with texture and motion cues that were in fact consistent ($d_t = d_m = 6.6$ cm) and serves as the *unperturbed* cue combination reference. The perturbation analysis considers the effects of varying one cue in a combination by comparing the *perturbed* psychometric functions ($d_t \neq d_m$) to the unperturbed function. A qualitative examination of the psychometric functions indicates that a cue perturbation shifts the psychometric function for the combined depth comparison in the same direction that the cue is shifted. It also appears that the shapes of the psychometric functions do not change much with the perturbation.

To make quantitative evaluations we fitted twoparameter cumulative probability distribution functions to the data for each mixed-cues surface. Maximum likelihood estimates (MLE) of parameters were computed using a downhill simplex method (Press, Flannery, Teukolsky & Vetterling, 1988, pp. 305–308). We fitted both normal and Weibull distributions and found that the normal distribution gave a better fit in four of the six sets of subject × condition, as indicated by χ^2 . We used the normal distribution parameters for further analysis. Of the six data sets, both sets for EBJ and one set (d_t fixed) for LTM resulted in normal distribution fits that were rejected by χ^2 ($\alpha = 0.01$).

The simplest linear model would consist of parallel psychometric functions, a perturbed single cue just shifting the combined depth estimate without affecting sensitivity, and this is consistent with qualitative observations. We tested this quantitatively by comparing the fit of a parallel set of psychometric functions to independently fitted curves. For each subject × condition there will be a different mean for each of the five distributions. Parallel functions will have a common variance while independent functions may each have different variances. With the MLE of the six parallel psychometric function parameters and the 10 independent psychometric function parameters, we can test the null hypothesis that the functions are parallel with the nested hypothesis test (Mood, Graybill & Boes, 1974, pp. 440–442). Comparing the likelihoods (L) of the two fits as

$$2\log_{e}\left(\frac{L_{\text{independent}}}{L_{\text{parallel}}}\right) \tag{7}$$

is distributed as χ^2 with 10-6=4 degrees of freedom. The likelihood that the data were sampled from normal distributions of equal variances was not significantly different from the likelihood that the data were sampled from distributions of different variances for four of six sets of subject × condition. At $\alpha = 0.01$, the parallel psychometric functions model was rejected for MJY and EBJ when d_t was fixed.

For a model with parameters $\theta_1, \theta_2, \ldots, \theta_n$, the joint distribution of MLE $(\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_n)$ of the parameters is asymptotically Gaussian. The expected value of the estimate $\hat{\theta}_i$ is asymptotically θ_i . For an $n \times n$ matrix M, the variance of $\hat{\theta}_i$ is $\sigma_{\hat{\theta}_i}^2$ where

$$\sigma_{\theta_i}^2 = -(M^{-1})_{ii}, \quad M_{jk} = \frac{\partial^2 (\log_e L)}{\partial \theta_i \partial \theta_k}$$
(8)

(Mood et al., 1974, pp. 358-362; Kendall & Stuart, 1979, pp. 59-64). Our analysis uses the fitted psychometric functions' parameters assuming the estimates are sampled from normal distributions with standard deviations given by equation (8). Based upon the partial success of the parallel model in the nested hypothesis test and the absence of statistically significant differences between MLE parameters for parallel and independent psychometric functions, we used the more compact parallel model for the subsequent analysis. The mean of the standard deviation parameters of the fitted normal distributions across all subjects and psychometric functions was 1.5 cm. This establishes the size of single cue perturbations (1.1 and 2.2 cm) to be about 0.75 and 1.5 standard units. Cue perturbations were then on the order of one jnd, confirming that the independent variables were manipulated within a range appropriate to the perturbation analysis.

The MLE of the mean of a cumulative normal distribution function corresponds to the psychometric function's 50% level. The mean is then the point of subjective equality (PSE), the depth of a consistent-cues surface judged to extend further in depth on 50% of the comparisons with the mixed-cues surface generating the psychometric function. We take the PSE to be the perceived depth of a mixed-cues surface. Each psychometric function then yields one PSE as a function of the perturbed depth cue in the mixed-cues surface [equation (1)]. These functions are plotted in Fig. 4. The data points are marked with error bars which indicate 99% confidence intervals for the measurement based upon standard deviations computed using equation (8).



FIGURE 4. A plot of PSE as a function of the perturbed depth cue. Panels are arranged as in Fig. 3. Each psychometric curve in Fig. 3 corresponds to one point on this plot. The error bars represent 99% confidence intervals for each PSE. The lines are maximum likelihood regressions of PSE on the portrayed depth of the perturbed cue in the mixed-cues stimulus. The data are generally consistent with a linear model.

Straight lines were fitted to PSE measurements by MLE. We tested whether the data were consistent with a linear relationship between PSE and the perturbed cues by comparing the confidence intervals for each PSE to the distance between the PSE and the fitted line. As this entails performing five independent tests per fitted line, we set the significance level to $1 - (1 - \alpha)^{1/5}$ for each comparison. The overall significance level (probability of incorrectly rejecting the linear hypothesis for any of the five points associated with a line) was set to $\alpha = 0.01$. We could reject the linear hypothesis for only one of the six fits (subject LTM with d_m perturbed). Note however that the test of linearity is unavoidably weak because we have no prior hypothesis concerning the pattern of nonlinearity to be expected under the alternative hypothesis.

The slopes of the lines and the sums of the slopes are listed in Table 1. The confidence intervals are computed under the assumption that the estimates are normally distributed with variances derived from PSE variances given by equation (8). The sums of the weights are not significantly different from one ($\alpha = 0.01$), although this is a consequence of the choice of stimuli (see the

TABLE 1. Depth cue weights from Expt 1

	MJY	EBJ	LTM
α,	0.46 (0.39-0.53)	0.53 (0.44-0.62)	0.27 (0.17-0.37)
α _m	0.45 (0.38-0.52)	0.45 (0.36-0.54)	0.61 (0.51–0.71)
$\alpha_t + \alpha_m$	0.91 (0.81–1.01)	0.98 (0.85–1.11)	0.88 (0.74–1.02)

Ranges in parentheses are 99% confidence intervals for each estimate.

Discussion of Expt 2). The subjects' combined depth estimates appear to be weighted linear combinations of the depth estimates from the individual cues.

The question then arises as to what determines the values of the weights. It may seem surprising that two of the subjects weighted texture information as heavily as motion information. The common subjective impression with more complex stimuli is that motion is a stronger depth cue than texture gradient. For example, Braunstein (1968) found the motion cue to have twice as much weight as the texture cue in a study which measured perceived slant with stimuli which included dot density and motion gradients on inclined planes. Pilot studies indicated that motion cues could be sufficiently weakened to make them more comparable in influence to texture cues. We then set our baseline stimulus parameters so as to yield relatively equal texture and motion weights.

If these data are partitioned by the rotation amplitude of the mixed-cues surface into two sets, one for total rotations about the line of sight of 10-20 deg, another for larger rotations of 20-30 deg, we derive two different sets of weights (Table 2). The weight given to the motion cue is increased for the cases of greater rotation amplitude (although not significantly for subject EBJ at a one-tailed $\alpha = 0.01$). With greater rotations a larger set of 2D measurements is available for computing structure from motion. The availability of information relevant to a depth computation affects the weight given to the cue (Landy *et al.*, 1991b). Notice also that in all cases the weight of the texture cue decreases as the weight of the motion cues increases (although none of the differences is significant at a one-tailed $\alpha = 0.01$).

From the viewpoint of statistical decision theory, we would also expect the reliability of the information used in computing depth from a particular cue to influence the weight that a normative observer would give that cue in a linear combination. Although this appears to be consistent with the increase of α_m as the rotation extent is increased, we tested this more directly in Expt 2.

EXPERIMENT 2

In the first experiment we demonstrated that depth judgments combining texture and motion cues are linear combinations with weights that may depend upon the availability of the relevant information in the scene. In this experiment we explore the role of reliability in determining the weights by artificially reducing the reliability of the single cues. We measure the weights with the same paradigm as used in Expt 1. Assuming that the depth information in an orthographic projection of a curved surface is the gradient of compression of the texture elements (Cutting & Millard, 1984), we may reduce the reliability of the texture cue by increasing the variability of texture element compression in the view. Assuming that the gradient of texture element velocity provides the information used to determine the structure from motion, we may reduce the reliability of the motion

	-	-		
Rotation		MJY	EBJ	LTM
10-20 deg	α,	0.44 (0.35-0.53)	0.59 (0.47-0.71)	0.29 (0.15-0.43)
e	α	0.39 (0.29-0.49)	0.43 (0.30-0.56)	0.32 (0.17-0.47)
	$\alpha_1 + \alpha_m$	0.83 (0.69-0.97)	1.02 (0.85-1.19)	0.61 (0.41-0.81)
20-30 deg	α,	0.36 (0.26-0.46)	0.51 (0.38-0.64)	0.27 (0.13-0.41)
U	à	0.53 (0.43-0.63)	0.49 (0.36-0.62)	0.59 (0.45-0.73)
	$\alpha_{t} + \alpha_{m}$	0.89 (0.75-1.03)	1.00 (0.81-1.19)	0.86 (0.67-1.05)

TABLE 2. Depth cue weights from Expt 1 partitioned by rotation amplitude

Ranges in parentheses are 99% confidence intervals for each estimate.

cue by increasing the variability of the texture element velocities in the view.

Methods

Stimuli. The stimuli were produced in the same fashion as Expt 1. The consistent-cues surfaces were the same as those used in Expt 1. However, the mixed-cues surfaces were constructed with some modifications.

For conditions with reduced texture cue reliability, the textures of the mixed-cues surfaces were generated by volumes packed with ellipsoids rather than spheres as in Expt 1. All ellipsoids had semiaxes parallel to the cylinder's elliptic semiaxes and vertical axis. The ellipsoid semiaxis parallel to the cylinder's vertical axis was the same length as the radii of the spheres used in Expt 1. The other ellipsoid semiaxes were equal in length to each other (i.e. ellipsoids of revolution) although that length varied randomly from ellipsoid to ellipsoid. The lengths of these ellipsoid semiaxes orthogonal to the cylinder's vertical axis were normally distributed over ellipsoids with mean equal to the ellipsoid vertical semiaxis (i.e. length of spheres' radii in Expt 1) and standard deviation one-half of that. (Actually, the distributions were clipped at 2 SDs, so that lengths ranged from 0 to 20 pixels.) The gradient of texture element aspect ratios then has a coefficient of variation (CV_1) of 0.5. The geometry of the cylindrical surfaces used in the stimuli were otherwise the same as for Expt 1. Images of nontilted cylinders with added aspect ratio variability can be seen in Fig. 5.

For the conditions with reduced motion cue reliability, the textures were generated as in Expt 1. The geometry and motion of the mixed-cues surfaces were the same as in Expt 1 except for an additional manipulation of the



FIGURE 5. Examples of stimuli for different levels of texture cue reliability. The upper row shows textures from Expt 1. The lower row shows textures from Expt 2 with noise added to the texture cue (for a $CV_t = 0.5$). Each column shows a surface for a given level of d_t .



FIGURE 6. The psychometric functions for depth comparisons with the mixed-cues surfaces with reduced texture cue reliability $(CV_t = 0.5)$ from Expt 2.

position of individual texture elements. For each motion sequence frame, after the 3D position of each texture element was computed based upon the cylinder's rotation, each texture element on the mixed-cues surface was additionally displaced along the cylinder's surface. The displacement was in a direction parallel to the cylinder's vertical axis. The magnitude of the displacement was sampled from a normal distribution with mean 0 and standard deviation one-half of the texture element's 3D displacement with respect to the previous frame. This generates an optic flow variability roughly proportional to velocity for a CV_m of 0.5, relevant to a depth from motion computation. Despite the fact that the individual texture elements now jitter back and forth somewhat independently, the displays still look like rotating cylinders.

Procedure and subjects. The procedure for Expt 2 was identical to that of Expt 1. Each single cue reliability level evaluated (CV_t for two subjects, CV_m for one subject) required the same set of measurements as Expt 1. The same number of trials and blocks and the same levels of the perturbed cue in the mixed-cues surfaces for both cues were repeated for each cue reliability level evaluated. Subjects MJY and EBJ are the same observers as participated in the first experiment. MJY and EBJ provided the measurements for CV_t . MJY provided the measurements for CV_m .

Results and discussion

Figure 6 shows the psychometric functions for depth comparisons with the mixed-cues surfaces with increased texture element aspect ratio variability in the mixed-cues surfaces. If these psychometric functions are compared to those in Fig. 3 one will note the general effect of the perturbed cues is similar, the psychometric functions are



 $d_{\tau} = 6.6 \text{ cm}$

d₁= 6.6 cm

FIGURE 7. PSE as a function of the perturbed depth cue in mixed-cues surfaces with reduced texture cue reliability ($CV_t = 0.5$). Each point corresponds to a psychometric function in Fig. 6. Note that the slopes in the left-hand panels are less than the corresponding panels of Fig. 4, suggesting that the texture cue is given less weight.

shifted along the consistent-cues' depth axis in the direction of the perturbation. These curves are also roughly parallel. Psychometric functions with d_t varied for MJY and d_m varied for EBJ pass the nested hypothesis test for parallel psychometric functions ($\alpha = 0.01$).

It is important to note that the functions for the perturbed texture cue are more closely spaced for $CV_1 = 0.5$ than for Expt 1. This effect is more clearly seen in the slopes of the functions of PSE for single cue perturbations graphed in Fig. 7. As in Expt 1, the combined depth estimate is a linear combination of the two cues, as demonstrated by the linearity of PSE with single cue perturbation. Different from Expt 1 are the slopes of the comparable lines. Compare the functions graphed in Fig. 7 with the same subjects' functions in Fig. 4. The weights of the individual cues in the linear combination are listed in Table 3. These are to be compared with those in Table 1. In both cases, the weight given to the texture cue significantly decreased when the reliability of the texture information was reduced (one-tailed $\alpha = 0.01$). The weight given to the motion cue either did not change (EBJ) or increased in a nearly complementary fashion (MJY).

TABLE 3. Depth cue weights from Expt 2

Condition		МЈҮ	EBJ
$CV_{1} = 0.5$	α,	0.26 (0.19-0.33)	0.39 (0.29-0.49)
	αm	0.62 (0.55-0.69)	0.44 (0.34-0.54)
	$\alpha_1 + \alpha_m$	0.88 (0.78-0.98)	0.83 (0.69 0.97)
$CV_{m} = 0.5$	α,	0.30 (0.24-0.36)	
	αm	0.20 (0.14-0.26)	
	$\alpha_{t} + \alpha_{m}$	0.50 (0.41-0.59)	

Ranges in parentheses are 99% confidence intervals for each estimate.





FIGURE 8. The upper row shows the psychometric functions for depth comparisons with the mixed-cues surfaces with reduced motion cue reliability ($CV_m = 0.5$) from Expt 2. The lower row shows PSE as a function of the perturbed depth cue, each point corresponding to a psychometric function in the plot above. The slope in the lower-right-hand panel is less than the corresponding panel of Fig. 4, suggesting that the motion cue is given less weight.

The story is similar for MJY when the reliability of the motion cue is reduced. The upper row of Fig. 8 shows the psychometric functions for MJY's depth judgments for trials in which the mixed-cues surfaces had increased optic flow variability. The perturbed cue again shifts the psychometric function along the depth axis relative to the unperturbed central function $(d_{\rm t} = d_{\rm m})$. Compared to the psychometric functions in Fig. 3, however, the shifts are smaller. Note how closely spaced are the psychometric functions for the perturbed motion cue. This close spacing is reflected in the flatter slopes of PSE as a function of the perturbed depth cue as graphed in the lower row of Fig. 8. In this case, as in the previous, increasing the variability of the information in a single depth cue does not destroy linearity, but it decreases the weight given to the cue in the linearly combined depth estimate. Compare the weights of the cues for the higher quality displays of Expt 1 (Table 1) with those for reduced motion cue reliability for subject MJY (Table 3). The motion cue's weight is dramatically reduced with reduced motion cue reliability. Unexpectedly, the texture cue's weight is also reduced.

The data also demonstrate that the depth of a surface with lower quality cues appears different than that of a surface of the same distal shape with higher quality cues. When measuring the unperturbed psychometric functions of Expt 1, comparisons are made between surfaces that are essentially the same (different only in that the random textures are sampled independently for each surface). Thus, the PSE for this function must match the unperturbed cue, by definition. In other words, in Fig. 4 a mixed-cues surface with $d_t = d_m = 6.6$ cm must match a consistent-cues surface with the same portrayed depth values (up to binomial variability). If perceived depth is a linear combination, as described by equation (2), this forces the sum of weights to be one, since

$$d = \alpha_{\rm t} d + \alpha_{\rm m} d = (\alpha_{\rm t} + \alpha_{\rm m}) d. \tag{9}$$

In Expt 1 the weights must sum to one because of the choice of stimuli. This is not the case in Expt 1 when we examine the subset of the data for which the mixed-cues rotation amplitude is constrained and the consistentcues rotation amplitude is unconstrained (Table 2); in this case it is meaningful to check whether the cue weights sum to one.

In Expt 2 the unperturbed psychometric function is generated by comparisons between a mixed-cues surface and a consistent-cues surface that are rendered from the same solid shape model, but with cues of different quality. Thus, Expt 2 does not constrain the measured weights to sum to one. We predicted that the slope of PSE as a function of perturbed depth cue would change with cue quality but had no prediction of an effect on the PSE-intercept. Note that the lines graphed in Fig. 7 and in the lower row of Fig. 8 mainly pass below the point where the perturbed cue is equal to depth to the unperturbed cue (6.6 cm), whereas the lines in Fig. 4 pass through that point. This implies that the surface with lower quality cues appeared to have less depth than one of the same distal shape with higher quality cues.

This depth effect of cue quality may be due to the cue combination process. The linear combination rule of equation (2) may be made consistent with the data from the conditions with reduced cue reliability (including the 10-20 deg subset of Expt 1) in two ways. One model suggests that decreased cue reliability somehow decreases the depth computed from an individual cue (e.g. d_m is decreased by motion noise). Alternatively, lowered reliability may have no affect on the depth computed using a cue (other than to reduce the reliability of the depth estimate), but overall depth is reduced due to a change in weight (e.g. a reduction in α_m). This reduction in cue weight only implies a reduced overall depth if the other cue weight (α_t) is not increased in a compensatory fashion, resulting in $\alpha_t + \alpha_m < 1$. We refer to this property as subadditivity. In the $CV_t = 0.5$ condition of Expt 2 it is clear that the results are not consistent with merely a change in depth from texture (d_t) : although overall depth was decreased slightly by the unreliable texture cue, the measured weight of the motion cue increased in a nearly compensatory fashion. However, in the $CV_{\rm m} = 0.5$ condition the data do not allow us to distinguish these two explanations because overall depth did decrease. Below we outline several reasons why the sum of the weights might change with cue quality.

CONCLUSIONS

The results of these experiments suggest, first, that perceived depth is linearly dependent on the depth encoded in the individual cues, as evidenced by the fits of PSE as a function of the depth portrayed by the perturbed cue. Second, they suggest that weights change with changes in the reliability of the different types of depth cues. Unreliable cues are downweighted in favor of reliable cues.

A weighted linear combination rule has a number of implications for quantitative studies of depth modules. We develop one such implication next.

Let us assume that the depth estimate from texture is perturbed by additive error with mean zero, variance σ_t^2 , that the depth estimate from motion is perturbed by additive error with mean zero, variance σ_m^2 , and that the errors in the depth and texture estimates are independent. Then we can compute the variance of the combined estimate of depth (using elementary properties of covariance and independence) as:

$$\sigma_{\rm d}^2 = \alpha_{\rm t}^2 \sigma_{\rm t}^2 + \alpha_{\rm m}^2 \sigma_{\rm m}^2. \tag{10}$$

Thus, σ_d^2 is a function of α_t and α_m , the weights assigned to texture and motion, respectively. When α_t and α_m sum to 1, we may replace α_t by $1 - \alpha_m$ in the above equation, and equation (10) is then a function of α_m which achieves its minimum at

$$\alpha_{\rm m}^{\rm opt} = \frac{\sigma_{\rm t}^2}{\sigma_{\rm t}^2 + \sigma_{\rm m}^2}.$$
 (11)

This choice of α_m minimizes the variance of the combined depth estimate. Conversely, if we know this optimum value, then we can solve for the ratio between the variances, σ_m^2/σ_t^2 :

$$\frac{\sigma_m^2}{\sigma_t^2} = \frac{1}{\alpha_m^{\text{opt}}} - 1.$$
(12)

Therefore, if we assume that the observer's choice of α_m is optimal then we can, in principle, estimate the ratio of variances using equation (12), at least for those cases where the sum of the measured weights is (approximately) 1. These results are readily generalized to *n* cues. Again, the sum of the *n* weights must be 1.

Drawing conclusions based upon the assumption that weights sum to one requires a full inventory of depth cues in a stimulus, and we have evidence that this is more difficult than is ordinarily assumed. Subadditivity of the depth cue weights was particularly pronounced for displays of reduced quality (see Table 3 and the 10-20 deg subset of Table 2). The reduction of a cue's weight with a reduction in the cue's quality was not always accompanied by a complementary increase in the other cue's weight (sometimes there was even a decrease). We take this as an indication that other cues are used and reducing the quality of a controlled cue may cause the weights to be redistributed to the advantage of cues we have not accounted for. The plots of PSE are also shifted downwards (see Figs 7 and 8), suggesting that these unidentified cues indicate flatter surfaces than those simulated experimentally. These cues may correspond to whatever information is still present in the experimental situation to indicate that the stimuli are in fact on the nearly flat surface of a CRT. The lighting and arrangement of the experimental situation were designed to weaken or eliminate these cues. By this account, the effect of these cues is unobservable when the experimental cues are "strong", but then when we reduce the reliability of the experimental cues as in Expt 2, these cues are given corresponding larger weights.

The analysis of results generated by stimuli with weak depth information may require consideration of *ex*traneous cues for depth, present in the scene, that typically signal that the depth stimulus is a flat pattern on a CRT screen. While the difficulty this adds to cue combination studies is obvious, it should be noted that it is also relevant to single cue studies. If the perceived depth in a single cue experiment is a weighted linear combination of the depth cue (d_{exp}) and the cues of the flat screen ($d_{flat} = 0$), equation (2) becomes

$$d = \alpha_{\exp} d_{\exp} + \alpha_{\text{flat}} d_{\text{flat}} = (1 - \alpha_{\text{flat}}) d_{\exp}.$$
(13)

There are then two classes of depth effects possible between experimental conditions, those due to changes in the depth estimated by the module (d_{exp}) and those due to a redistribution of weights that varies the influence of the extraneous cues (α_{flat}) . Reducing the quality of the depth cue may reduce perceived depth by increasing α_{flat} .

We hypothesize that in experimental situations where there are a sufficient number of strong depth cues under experimental control to make any extraneous cues insignificant, reducing the quality (e.g. perceived reliability) of one cue should not result in reduced perceived depth. Rather, more weight is given to the remaining, more reliable cues. However, when there are few reliable cues in the display, then a reduction in single cue reliability will result in less perceived depth (that is, more weight given to the extraneous cues that typically signal "flatness"). To test this idea, we constructed a demonstration that employed a finding by Loomis and Eby (1988): perceived depth of KDE displays decreases with a decrease in the angle between the axis of rotation and the line of sight. This makes sense since in the limit of rotation about the line of sight, no information is available in the display for computing depth from motion. They used multidot displays which give only a weak texture cue (only dot density varies across the display). Thus, by our account their results do not necessarily indicate a decrease in the estimate of depth from motion, but rather a decrease in the weight applied to that estimate (and a corresponding increase in the weights of other cues which specify a flat display). Using the stimulus model from the experiments, we constructed stereograms of textured cylinders rotating about axes with various orientations relative to the line of sight. Perspective projection was employed. The depths of identical surfaces defined by perspective, texture compression, binocular disparity, and motion cues were compared and depth did not appear different for rotations about axes making angles with the line of sight

of 30, 45, or 90 deg (in the above experiments, the rotation axis was perpendicular to the line of sight). This demonstrates that the orientation of the rotation axis does not affect the perceived shape of an enriched depth stimulus (within the range tested), supporting our hypothesis.

Other observations were provided by the demonstration. Stationary cylinders or cylinders rotating about the line of sight appeared to have less depth than identical cylinders rotating about other axes. One explanation is that extraneous cues achieve significant weight in versions of our stimuli containing only disparity, perspective, and texture compression cues to depth. Another explanation is that the motion cue affects calibration of stereopsis, a cue that in isolation is miscalibrated (Johnston, 1991; Johnston, Cumming & Landy, 1992). Either way, this also demonstrates that for our displays with rich depth information, independence of perceived depth from rotation axis orientation is not due to a dominating disparity cue. Instead, the motion cue must have a weight that is transferred to the extraneous cues when motion is eliminated or irrelevant for depth estimation. The monocular percept was that depth decreased as the angle between the rotation axis and the line of sight decreased, as Loomis and Eby found. Thus, extraneous cues also achieve significant weight when our stimuli contain only perspective, texture, and reduced quality motion cues.

If the subadditivity of weights observed is, in fact, due to extraneous cues, then it could be eliminated by an experimental apparatus that provides a sufficiently compelling simulation of depth stimuli. The sum of the weights of the controlled cues then serves as an index of depth display quality. In some of the conditions of our experiments, this index was rather low. While this does not weaken our demonstration of a linear combination rule for depth cues, it complicates the analysis of the weight assignment process.

Interpretations of psychophysical measurements of depth cue combinations may be more complex when the sum of the weights is not known. For example, perceived depth increases with the number of cues included in the stimulus (Bruno & Cutting, 1988; Bülthoff & Mallot, 1988). This phenomenon may be due to redistributions of weights whereby extraneous cues to flatness are offset by additional depth cues. This is in contrast with the unnormalized additivity suggested by Bruno and Cutting, or regularization mechanism suggested by Bülthoff and Mallot. The extraneous cue hypothesis can be rejected when the weights of the controlled cues sum to one in each condition.* We may then consider explanations of this phenomenon that attribute it to an inherent response of the visual system to inconsistent or impoverished visual data, rather than a response to imperfect stimuli (e.g. actual stimulus cues to flatness). One alternative explanation is that when the visual system is

presented with unreliable or impoverished depth cues, it responds by producing flattened or "smoothed" estimates of surfaces in the scene. This oversmoothing hypothesis is also consistent with these results. This phenomenon may be due to a tendency for the visual system to perceive stimuli as flat [the equidistance tendency (Gogel, 1965)]. It may also arise due to a decrease of the "smoothing" term in a regularization scheme as more information becomes available (Bülthoff and Mallot, 1988).

Consideration of either the extraneous cue hypothesis or the oversmoothing hypothesis suggests that it may be difficult to study the integration of depth modules with stimuli that have impoverished depth information, stimuli that are made unreliable, or with depth cues that signal markedly discrepant depths at the same location in a scene. The visual system may react to unreliability or inconsistency *per se* by altering the depth combination rule. The perturbation analysis methods and accompanying psychophysical procedures are designed to circumvent such problems.

In summary, the analysis of depth vision with cues strongly in conflict has proven to be difficult. A full description of the depth combination rules for inconsistent and impoverished stimuli seems likely to resemble a microcosm of cognitive processing: elements of memory, learning, reasoning and heuristic strategy may dominate. In contrast, we suggest that under conditions where reliable and consistent depth cues are presented, and the experimental procedure does not itself reduce reliability or render consistent cues inconsistent, the depth combination rule may prove to be both simple and normative.

REFERENCES

- Aloimonos, J. & Shulman, D. (1989). Integration of visual modules: An extension of the Marr paradigm. New York: Academic Press.
- Braunstein, M. L. (1968). Motion and texture as sources of slant information. Journal of Experimental Psychology, 78, 247-253.
- Bruno, N. & Cutting, J. E. (1988). Minimodularity and the perception of layout. Journal of Experimental Psychology: General, 117, 161-170.
- Bülthoff, H. H. & Mallot, H. A. (1988). Integration of depth modules: Stereo and shading. Journal of the Optical Society of America A, 5, 1749–1758.
- Clark, J. J. & Yuille, A. L. (1990). Data fusion for sensory information processing systems. Boston, Mass.: Kluwer.
- Cutting, J. E. & Millard, R. T. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General*, 113, 198–216.
- Dosher, B. A., Sperling, G. & Wurst, S. A. (1986). Tradeoffs between stereopsis and proximity luminance covariance as determinants of perceived 3D structure. *Vision Research*, 26, 973–990.
- Epstein, W. (1968). Modification of the disparity-depth relationship as a result of exposure to conflicting cues. *American Journal of Psychol*ogy, 81, 189-197.
- Gibson, J. J. (1950). The perception of visual surfaces. American Journal of Psychology, 63, 367-384.
- Gogel, W. G. (1965). Equidistance tendency and its consequences. Psychological Bulletin, 64, 153-163.
- Johnston, E. B. (1991). Systematic distortions of shape from stereopsis. Vision Research, 31, 1351–1360.
- Johnston, E. B., Cumming, B. G. & Landy, M. S. (1992). Integration of stereopsis and motion shape cues. Mathematical Studies in Perception and Cognition, 92-1, New York University, New York.

^{*}A weaker rejection of the extraneous cue hypothesis might be possible when the sum of the weights of the controlled cues is constant across conditions.

- Johnston, E. B., Cumming, B. G. & Parker, A. J. (1993). Integration of depth modules: Stereopsis and texture. Vision Research, 33, 813–826.
- Julesz, B. (1971). Foundations of Cyclopean perception. Chicago, Ill.: University of Chicago Press.
- Kendall, M. K. & Stuart, A. (1979). The advanced theory of statistics: Vol. 2. Inference and relationship (4th edn). New York: Macmillan.
- Landy, M. S., Maloney, L. T. & Young, M. J. (1991a). Psychophysical estimation of the human depth combination rule. Sensor Fusion III: 3-D Perception and Recognition, Proceedings of the SPIE, 1383, 247-254.
- Landy, M. S., Maloney, L. T., Johnston, E. B. & Young, M. J. (1991b). In defense of weak fusion: Measurement and modeling of depth cue combination. Mathematical Studies in Perception and Cognition, 91-3, New York University, New York.
- Loomis, J. M. & Eby, D. W. (1988). Perceiving structure from motion: Failure of shape constancy. International Conference on Computer Vision.
- Maloney, L. T. & Landy, M. S. (1989). A statistical framework for robust fusion of depth information. Visual Communications and Image Processing IV. Proceedings of the SPIE, 1199, 1154-1163.
- Marr, D. (1982). Vision: A computational investigation into the human representation and the processing of visual information. San Francisco, Calif.: Freeman.
- Mayhew, J. E. W. & Frisby, J. P. (1981). Psychophysical and computational studies towards a theory of human stereopsis. Artificial Intelligence, 17, 349-386.
- Mood, A. M., Graybill, F. A. & Boes, D. C. (1974). Introduction to the theory of statistics (3rd edn). New York: McGraw-Hill.
- Ono, M. E., Rivest, J. & 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.
- Peachey, D. R. (1985). Solid texturing of complex surfaces. Computer

Graphics, 19, 279-286.

- Press, W. H., Flannery, B. P., Teukolsky, S. A. & Vetterling, W. T. (1988). Numerical recipes in C: The art of scientific computing. Cambridge: Cambridge University Press.
- Richards, W. (1985). Structure from stereo and motion. Journal of the Optical Society of America A, 2, 343–349.
- Sperling, G., Landy, M. S., Dosher, B. A. & Perkins, M. E. (1989). Kinetic depth effect and identification of shape. Journal of Experimental Psychology: Human Perception and Performance, 15, 826-840.
- Todd, J. T. & 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.
- Wallach, H. & Karsh, E. B. (1963). The modification of stereoscopic depth-perception and the kinetic depth effect. *American Journal of Psychology*, 76, 429–435.
- Wallach, H. & O'Connell, D. N. (1953). The kinetic depth effect. Journal of Experimental Psychology, 45, 205–217.
- Young, M. J., Landy, M. S. & Maloney, L. T. (1991). Depth from texture and motion. *Investigative Ophthalmology and Visual Science* (Suppl.), 32, 1180.

Acknowledgements—The work described in this paper was supported in part by grant EY08266 from the National Eye Institute, NIH. In addition, Michael Landy was supported by an NRC-NASA Ames Research Associateship. We would like to thank Elizabeth Johnston for her helpful comments, two anonymous reviewers for their comments on a previous draft, and Robert Picardi for technical assistance. Portions of this work have been presented at the annual meetings of the Association for Research on Vision and Ophthalmology (Sarasota, Fla, May 1991) and the SPIE (Boston, Mass., November 1990) and published in abstract or proceedings form (Landy *et al.*, 1991a; Young, Landy & Maloney, 1991).