



Independent Encoding of Surface Orientation and Surface Curvature

ALAN JOHNSTON,* PETER J. PASSMORE*f

Received 27 May 1993; in revised form 13 December 1993

Marr [(1982) *Vision*, San Francisco, Calif.: Freeman] proposed that we represent surface geometry in terms of a viewer-centred description of surface orientation and distance. This description is computed by a range of independent processing systems which take as input particular kinds of information present in images, like surface texture, shading, retinal disparity and motion parallax. The outputs of these modules are integrated in order to provide a unitary representation of the layout of visible surfaces. Higher order properties of surface geometry, like surface curvature, might be computed from this symbolic representation or might be encoded independently from the visual information available at the retinae. We measured surface slant and surface curvature discrimination thresholds for surface patches defined by shading, texture and retinal disparity as a function of the elevation of the illumination. We found that observers judgements about the curvature of local surface patches were too precise to be based on a symbolic representation of surface orientation and we conclude that surface curvature is computed directly from depth cues present in the retinal images.

Shading Texture Stereopsis Surface representation Shape

INTRODUCTION

There are many ways to describe the shape of objects and environments. The simplest description which captures surface geometry is the range map, in which the distance to a surface is assigned for each direction in the visual field. Between a range map and global characterisations of shape lie a variety of descriptors which could be used to form a rich representation of surface layout (Barrow & Tenenbaum, 1978; Brady, Ponce, Yuille & Asada, 1985; Koenderink, 1990; Marr, 1978, 1982). Possible primitives include surface normals and gradients, mean and Gaussian curvature, the relative magnitude of the principle curvatures (Koenderink, 1990) and higher order descriptors like furrow, hump and ridge, which characterize the embedding of a particular surface type in a region which has a different classification (Koenderink & van Doorn, 1980). General descriptions can also be assigned for simple objects such as spheres, cones, tori and ellipsoids or surface qualities like undulating or cratered. However, there is little psychophysical evidence available on the validity of these shape descriptors for human vision.

Horn (1975) and Marr (1978) emphasized the surface orientation map as a means of representing surface shape. Two parameters are required to specify

surface orientation. The most natural parameterization is in terms of slant, the angle that the surface makes with the line of sight, and tilt, the direction of surface slant (Stevens, 1983). Psychophysical studies of surface orientation perception from shading and texture cues have generally found that observers underestimate surface slant (Braunstein, 1976; Gibson, 1950; Perrone, 1980). Surfaces appear to be flatter than they ought to. Mingolla and Todd (1986) asked subjects to report the slant and tilt of the surface normal at a point on a computer generated surface which was marked by a small cross. They found the errors involved in this task were high—mean errors were typically > 13 deg. Koenderink, van Doorn and Kappers (1992) confirmed that subjects judgements of surface orientation were highly variable using a perceptual conformity task in which a probe stimulus was aligned with the surface normal and tangent plane, although judgements of surface tilt were more consistent than judgements of surface slant.

Findings like these have led to the view that shading provides a poor cue on which to base a precise description of surface geometry. Todd and Reichel (1989) have suggested that, in the case of shading and texture cues, a metric approach to representation based on local mappings of depth and orientation has little psychological or perceptual relevance. The imprecision in subjects' judgements about surface geometry is taken as evidence that the representation of surface layout is, at best, coarse grained, and they argued in favour of an ordinal

*Department of Psychology, University College London, London WC1E 6BT, England.

fPresent address: School of Computing, University of North London, 166-220 Holloway Road, London N7 8DB, England.

representation of depth information. An ordinal scale would allow the signs of depth differences to be represented. A nominal representation based on the classification of surface patches as elliptic (principle curvatures have the same signs), hyperbolic (principle curvatures have opposite signs), or parabolic (one of the principle curvatures equals zero) might also be considered. However, Erens, Kappers and Koenderink (1993) report subjects have considerable difficulty in classifying quadratic surfaces on the basis of shading alone. Context may be important. Koenderink and van Doom (1980) point out that elliptical and hyperbolic regions on the surface are separated by parabolic lines. In the generic case, parabolic lines do not overlap, allowing a surface classification based on embeddings of one surface type in another which form furrows or ridges.

In contrast to the substantial errors in surface orientation judgements reported by Mingolla and Todd (1986), and Weber fractions of around 0.4 reported by Reichel and Todd (1991) for a relative depth discrimination task, Johnston, Passmore and Morgan (1991) and Johnston and Passmore (1994) report Weber fractions of close to 0.1 for a curvature discrimination task in which shading provided the only cue to shape, demonstrating that shading provides precise information about surface curvature.

There are at least two possible ways in which surface curvature might be encoded by the human visual system. Curvature might be recovered by first encoding surface orientation and surface height then computing mean curvature by operating on this symbolic representation (Carman & Welch, 1992) (see Fig. 1). This method predicts a sequential dependency in shape processing. Stimulus factors affecting surface orientation discrimination should affect surface curvature discrimination in a similar fashion. However, it is difficult to see how precise discrimination of curvature could be achieved, given the error variance found in surface height and orientation judgement tasks, if this sequential symbolic strategy was implemented in the human visual system. The alternative, that curvature is encoded directly from the image intensities, appears the more likely scheme. Rogers and Cagenello (1989) report that subjects can make precise judgements about surface curvature in stereoscopic

displays. In a task in which subjects had to judge whether a cylindrical surface was convex or concave the disparity range of their display at threshold was one-third of that required to detect a change in surface slant over the same spatial extent. In order to investigate the functional architecture of shape analysis in more detail we compared discrimination thresholds for changes in surface curvature and surface orientation in a surface alignment task for spherical patches defined by combinations of shading, texture and binocular disparity cues.

GENERAL METHODS

Subjects

The authors served as subjects in the first three experiments. Both subjects had well corrected vision and were well practised in these tasks. In the final experiment we used six unpractised subjects who were naive to the purpose of the experiment. Two of these subjects were familiar with psychophysical procedures.

Stimulus generation and display

A computer generated image of a sphere was constructed using ray casting techniques (Foley, van Dam, Feiner & Hughes, 1990) and rendered using the Phong illumination model (Phong, 1975). Full details of the stimulus generation are given in Johnston and Passmore (1994). The lighting parameters were chosen to model the effect of a point source on a perfect diffuse reflector or Lambertian surface. The ambient illumination was set at 0.1 and the direct illumination was set at 0.7. The light source was always 100 cm from the centre of the sphere.

Texture could be mapped onto the sphere using the equidistant azimuthal mapping. The texture map provides the albedo value for any point on the sphere. We used bilinear grey level interpolation to provide values at intermediate points in the map. The textured map was generated by band-limiting a random 256 x 256 grey level image using an ideal filter with a 1 octave bandwidth and a central frequency of 57.6 c/image. The band-limited noise texture was used to reduce aliasing and to support the use of bilinear grey level interpolation in the indexing

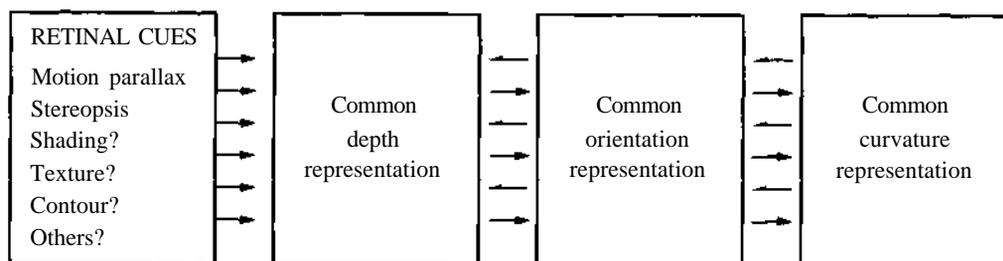


FIGURE 1. The symbolic pipeline architecture. It is possible that surface orientation and curvature could be computed from range data, without reference to the retinal images, by a process akin to differentiation. Similarly, lower order descriptions could be computed from higher order descriptions from a process akin to integration. Adapted with permission from Carman and Welch (1992).

of the texture map. This technique provides a regular isotropic texture whose spatial scale can be easily manipulated. It also ensures that Gibson's invariant, that equal amounts of texture map onto equal amounts of surface, holds true, at least for radial directions on the sphere. This is not the case for sculpted solid textures because the extent of a surface element that is painted a particular colour depends upon the angle at which the voxels are sectioned. In the texture mapping technique distortion of texture occurs in the tangential direction. Distortion is relatively slight for the region of the sphere used in these experiments (Johnston & Passmore, 1994).

For the stereoscopic displays we used a single mirror stereoscope. The image viewed by the left eye was generated by ray tracing from the viewpoint. The image viewed by the right eye was ray traced from the viewpoint and then reflected around the vertical axis. The right eye image was drawn to the left of the binocular display and was brought into stereoscopic alignment by rotating a vertical front silvered mirror placed close to the right eye. Matt black card was used to restrict the view to a single binocular display. The front silvered mirror was in fact one face of a dove prism. The advantages of this method are that a stereoscopic display can be achieved with a single monitor without use of multiple mirrors, which need careful alignment, or the use of polaroid filters to select images, which leads to reductions in image brightness. Rotation of the single mirror allows control over vergence and rotation of the dove prism allows control over the degree of cyclotorsion. We found that minor distortions of the shape of the object caused by substituting a translation of the right eye image in the plane of the display for a rotation could be corrected by slight rotations of the display or the head.

Images were displayed with eight bit precision. The stimuli were presented on a 16 in. Sony Trinitron monitor screen under the control of a SUN Sparcstation II. In order to linearize the display the luminance of the display was measured using a micro-photometer and appropriate values were placed in the display lookup table (the colourmap).

Procedure

Subjects were asked to make judgements about either the curvature or the slant of a test patch defined on the computer generated sphere. The sphere, which had a diameter of 7.5 cm, was modelled as having its origin in the plane of the display screen—the *xy*-plane. The screen was viewed from a distance of 75 cm. The test patch was separated from the main body of the sphere by an annulus of uniform brightness. The diameter of the patch was 20% and the width of the annulus was 1.6% of the diameter of the sphere. The luminance value of the annulus was set to the luminance calculated for the central pixel of the patch prior to experimental manipulation. The annulus served to prevent subjects from using the emergence of a brightness transition at the boundary of the patch as a cue. In addition the albedo of the test patch was varied by $\pm 10\%$ over trials, which perturbs the average

brightness of the patch, local brightness gradients, and the brightness range. In the Phong model albedo jitter does not affect the Michelson contrast. The texture in the test patch was rotated to a particular orientation which was selected at random for each trial.

In some of the experiments the slant of the patch was varied systematically and observers were asked to decide whether the patch was slanted to the left or right. In other experiments the curvature of the patch was varied and subjects were asked to decide whether the patch was more curved or less curved than the main body of the sphere. We chose an alignment task to avoid the use of a reference which was similar to the test stimulus so that subject would not be able to perform the task on the basis of detecting a difference between the brightness patterns in test and standard patches. In control experiments (Johnston & Passmore, 1994), which compared thresholds for positive and photographic negative versions of the stimuli, we found curvature discrimination thresholds were higher for photographic negative demonstrating that subjects were using the geometric cue in these tasks.

Thresholds were measured using an adaptive method of constant stimuli: Adaptive Probit Estimation (Watt & Andrews, 1981). Threshold is defined as the standard deviation of the error distribution and corresponds to the 84% point on the psychometric function. Thresholds were based on 64 individual trials. Each data point is the r.m.s. of four different threshold determinations. The standard deviation of the four individual thresholds provides a measure of dispersion.

EXPERIMENT 1: EFFECT OF LIGHT SOURCE POSITION FOR SHADED OBJECTS

In a previous study (Johnston & Passmore, 1994) we found that increasing the elevation of the illumination decreased curvature thresholds and increased slant discrimination thresholds for surface patches defined by shading. The first experiment attempted to replicate this effect using test patches which were illuminated from the same light source direction to allow a direct comparison. Stimuli were viewed monocularly using the dominant eye. Surface curvature is defined as the inverse of the radius of the spherical test patch. Discrimination thresholds for slant and curvature, plotted as a function of the elevation of the illumination above the line of sight, are shown in Fig. 3. Moving the light source towards the line of sight reduced slant discrimination thresholds but increased curvature discrimination thresholds. Figure 2(c,d) shows the change in the surface normals at threshold for the 65 deg illumination conditions. The needles for the slant task are scaled down by a factor of 10. In this lighting condition, the curvature threshold was around 0.01 cm^{-1} and the maximum change in the orientation of a surface normal at threshold was 0.74 deg. The average change in the orientation of a surface normal at threshold in the slant discrimination task was 7.6 deg (range = 7.5-7.8 deg). Thus, the discrimination thresholds found for surface curvature cannot be explained on the basis of

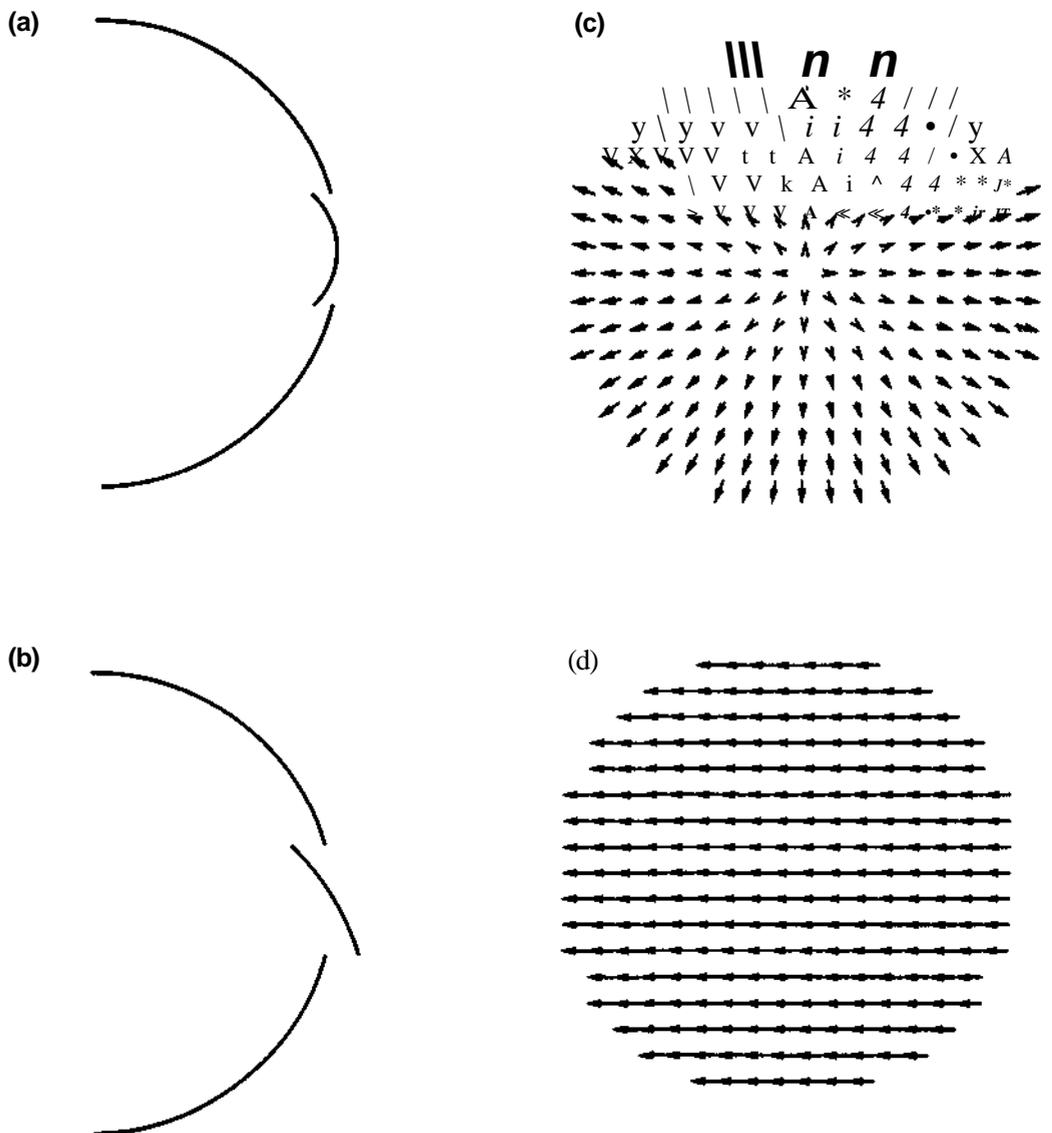


FIGURE 2. A sphere was rendered using ray tracing techniques and displayed on the Sony Trinitron monitor of a SUN Sparcstation II. The curvature (a) and orientation (b) of the central patch, shown here in horizontal cross-section, could be manipulated independently. The test patch was separated from the main body of the sphere by an annulus which was set to the brightest of the central point on the patch, (c) The changes in the surface normals in the curvature and (d) slant discrimination tasks. The length of the needles indicate the magnitude of the change at threshold for the monocular shaded images. The length of the needles in (d) are scaled down by a factor of 10. The orientation indicates the direction of the change.

detecting changes in the pattern of the surface normals in the patch.

EXPERIMENT 2: THE INTRODUCTION OF A BINOCULAR STEREOSCOPIC DEPTH CUE.

Shaded images cannot provide direct information about depth or distance and therefore, even if we accept the symbolic pipeline model described above, it is perhaps not surprising that we find a lack of correlation between the effects of illumination direction on slant thresholds and curvature thresholds in Expt 1. Binocular cues, however, provide information about surface range and orientation and the addition of these cues might allow evidence of sequential dependencies in the computation of surface curvature to emerge. In the next experiment the

shaded images were displayed stereoscopically. Figure 4 shows discrimination thresholds for the slant and curvature tasks plotted as a function of illuminant elevation as before. Again, moving the light source towards the line of sight reduced slant discrimination thresholds but increased curvature discrimination thresholds. All thresholds were lower using the stereoscopic display demonstrating that subjects were using the binocular cue. In addition, the slope of the function relating curvature discrimination threshold to light source elevation was reduced. Again curvature thresholds cannot be predicted on the basis of thresholds for surface slant. At threshold in the 65 deg illumination conditions the average change in the normals in the slant task exceeds the maximum change in the normals in the curvature task by a factor of 10.

EXPERIMENT 3: THE INTRODUCTION OF SURFACE TEXTURE.

It is generally considered that shaded displays provide a less favourable basis for stereo computation than textured displays (Blake, Zisserman & Knowles, 1985; Bülthoff & Mallot, 1988, 1990) so in the third experiment we introduced surface texture by mapping a band-limited random-noise texture onto the surface of the sphere. The presence of surface texture led only to reductions in slant discrimination thresholds and increases in curvature discrimination thresholds (Fig. 5). Slant thresholds are now independent of light source elevation but curvature thresholds rise more steeply as the elevation of the light source is reduced than in Expt 2. Again, for the 65 deg conditions, slant thresholds are too high, by a factor of 3.5, to predict performance in the curvature task. It is unlikely that subjects are using relative depth measures at the boundaries of the patch. For the 65 deg conditions the maximum depth change at threshold is just 0.02 mm in the curvature task and 0.11 mm in the slant task.

The decrease in slant discrimination threshold as the light source vector approaches the z-axis is in marked

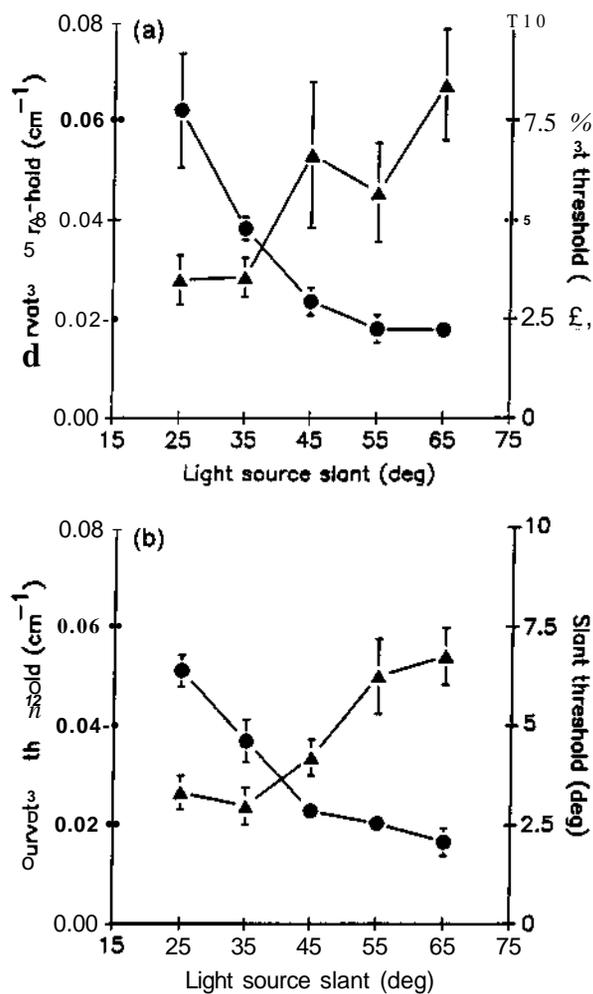


FIGURE 3. (a) Discrimination thresholds for the curvature (#) and slant (A) tasks as a function of the elevation of the light source above the line of sight for monocular viewing conditions. Curvature is expressed as the inverse of the radius of the spherical patch. Subject AJ. (b) Subject PP.

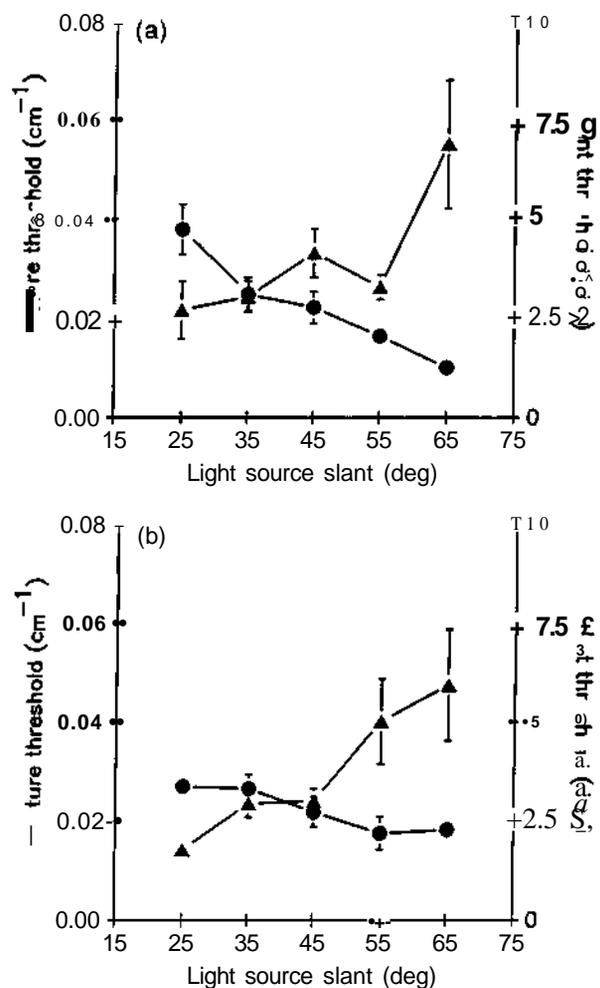


FIGURE 4. (a) Curvature (#) and slant (A) discrimination thresholds as a function of light source elevation for stereoscopically presented spheres. Subject AJ. (b) Subject PP.

contrast to the data for the curvature discrimination task where the reverse is the case. This argues strongly against the sequential symbolic model of shape encoding outlined earlier in which, it was suggested, surface curvature could be computed from a representation of surface orientation, and in favour of the proposal that curvature is encoded directly from the image values.

EXPERIMENT 4: THE EFFECTS OF ADDING TEXTURE ON SLANT AND CURVATURE DISCRIMINATION

Both the low Weber fractions for the curvature task and the finding that adding texture can increase discrimination thresholds for stereoscopic displays are surprising. However, our earlier experiments involved careful measurement with practised subjects and it is possible that the effects described above reflected familiarity with the task. In Expt 4 we compared curvature and slant discrimination thresholds for stereoscopic displays of shaded spheres with uniform or textured surfaces. In this case we used unpractised subjects who were naive to the purposes of the experiment. In the curvature task the

spheres were rendered using a light source set at 25 deg to the line of sight. For the slant task the light source was set at 65 deg elevation. Two of the subjects had some prior experience of psychophysical procedures but were not practiced on these tasks. The other four subjects were given 10-20 practice trials before data collection. For one subject we were not able to achieve a satisfactory fit to the psychometric function in one of the curvature conditions and only data for the slant task are included in the analysis. The naive subjects thresholds were based on 64 trials. The subjects completed the four conditions in a pseudo-random order. The data from the unpracticed subjects confirmed the earlier findings that curvature thresholds increase in the presence of the band-limited noise texture (Fig. 6; $I = 6.06$; d.f. = 4; $P < 0.05$). Thresholds are remarkably low and compare well with those of the practiced subject which are based on the average of four separate threshold determinations. The finding that the presence of texture decreases thresholds in the slant task was also confirmed (Fig. 6; $f = 2.74$; d.f. = 5; $P < 0.05$), however in this case thresholds varied greatly across subjects and average threshold values are considerably higher than those for the practiced subjects.

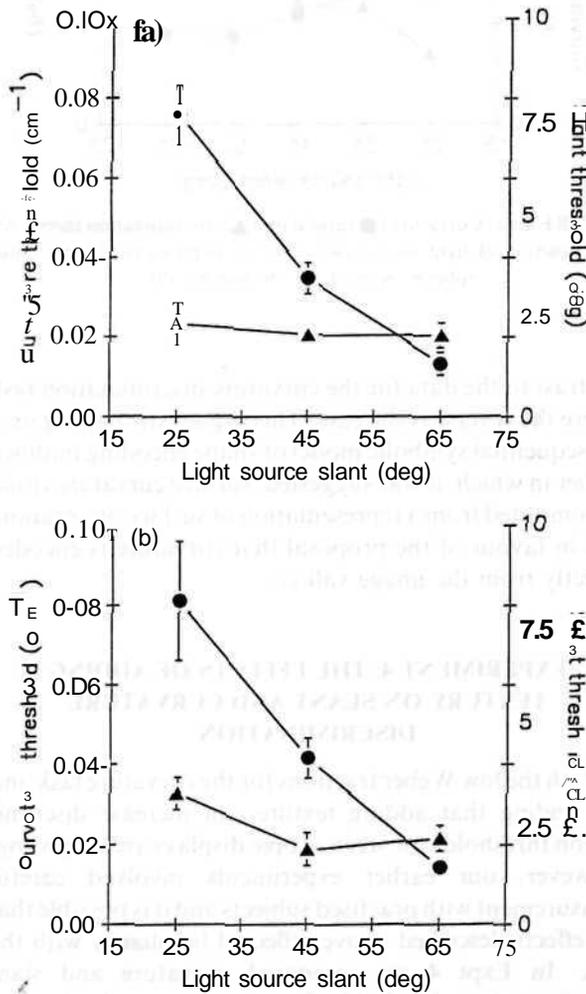


FIGURE 5. (a) Curvature (•) and slant (A) discrimination thresholds for stereoscopically presented spheres which had a band-limited random-dot surface texture. Texture contrast = 0.5. Subject AJ. (b) Subject PP.

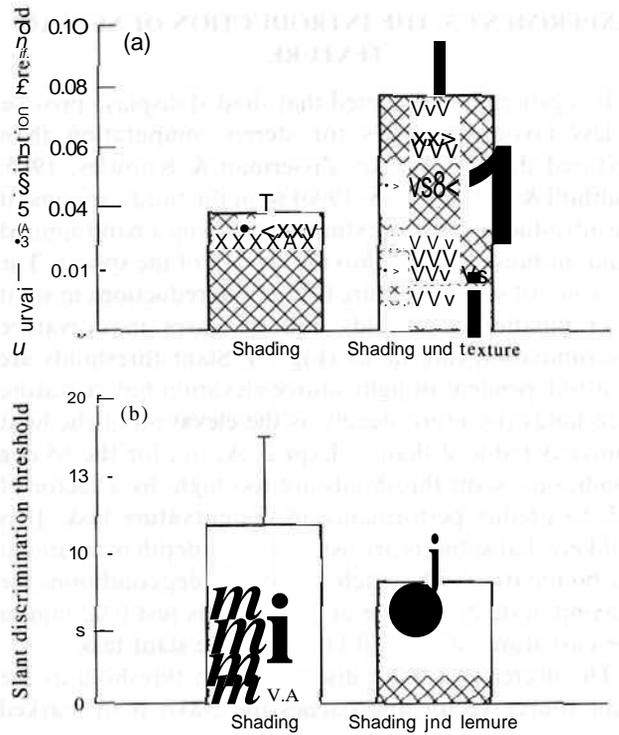


FIGURE 6. Average discrimination thresholds for the curvature discrimination task and the slant discrimination task for naive unpracticed subjects. The data are shown for two rendering conditions, uniform albedo and textured (50% contrast). In the curvature task the spheres were rendered using a light source set at 25 deg to the line of sight. For the slant task the light source was set at 65 deg elevation. The bar indicates 1 SE.

The relative ease of the curvature task for naive subjects and the observation that the presence of texture increases the difficulty of the curvature task while improving performance in the slant task is further evidence of the independence of the mechanisms for the recovery of information about surface slant and surface curvature.

DISCUSSION

Stevens (1984) has argued that the detection of changes in higher order properties of a surface, like surface curvature, may be mediated by detecting changes in the topography of surface normals. There is no need, he would argue, for an explicit representation of surface curvature in order to explain our ability to detect changes in surface shape. However, in our experiments, moving the light source towards the line of sight reduces slant discrimination thresholds but increases curvature discrimination thresholds. Thus, the discrimination thresholds found for surface curvature cannot be explained on the basis of detecting changes in the pattern of the surface normals in the patch. In addition to the interactions found between the light source elevation variable and the task, we also found that adding texture increased curvature thresholds but decreased slant discrimination thresholds. These observations also provide strong evidence against the idea that surface curvature is encoded from an explicit symbolic representation of surface orientation and in favour of the

proposal that surface shape and orientation are encoded in parallel from depth cues in the retinal images. Similar arguments apply to the question of whether subjects based their responses on depth or range information. If performance was limited by the precision with which distance was recovered we would expect changes in light source direction to have similar effects in the two tasks.

Adding texture improved subjects judgements of orientation. This would be predicted by any algorithm for the recovery of slant which depended upon measuring intensity variations or detecting edge features in the two retinal images. Textured objects provide well defined image features and substantive intensity gradients. However, the improvement in performance in the curvature task due to the introduction of disparate shading is interesting. Using a task in which subjects had to set a disparity probe to reflect the perceived depth of a point on ellipsoids defined by texture, shading and binocular disparity Bülthoff and Mallot (1990) demonstrated that psychophysical observers can make use of disparate shading to recover depth although they report that the presence of edges such as those generated by surface texture improved the accuracy of depth judgements. They also showed that ellipsoids which do not give rise to zero-crossings in Laplacian filtered images still give a good impression of stereoscopic depth. This is surprising given the emphasis on edge-based stereopsis in the computer vision literature (Frisby & Pollard, 1991). The texture used in the present experiments was a band-limited random-noise pattern. This texture is isotropic, cyclic and would give rise to a regular pattern of zero-crossings in Laplacian filtered images yet we find that adding surface texture reduces performance in the binocular curvature discrimination task. It is difficult to see how an edge-based stereo approach which leads to a description of surface distance could account for this effect. Introducing well defined image features should improve performance on any three-dimensional-shape task.

The pattern of the data suggests separate mechanisms for the encoding of surface curvature and for the encoding of depth and/or orientation. There may be some functional basis for dealing separately with these measures of surface geometry since depth and surface orientation provide information about the relationship between surfaces and some spatial reference frame whereas surface curvature [and particularly Koenderink's local shape parameter, S , the relative magnitudes of the values of the two principle curvatures in a surface patch (Koenderink, 1990)] describes properties of surfaces which are intrinsic and therefore invariant under changes in the reference frame (Johnston, 1992).

The idea that there may be different strategies for the recovery of curvature and slant has been considered by Rogers and Cagenello (1989) who demonstrated that curvature judgements in stereoscopic displays did not depend upon prior encoding of point disparities and they suggested a special mechanism for the computation of disparity curvature. They proposed that a biologically plausible method of estimating disparity curvature might

involve comparing the curvatures of planar line segments in binocular images. Cagenello and Rogers (1993) advocate a different mechanism for the recovery of surface slant based on the comparison of orientation differences in the two retinal images. However, this approach would not be able to provide an account for the decrease in thresholds found with the introduction of disparate shading or why texture was found to interfere with shape from disparate shading.

Although it is not clear why the presence of texture increases curvature discrimination thresholds the pattern of data is consistent with the notion of a separate mechanism for the computation of curvature from disparate shading. It is likely that any shape from shading mechanism would incorporate the assumption that brightness changes in the image result from changes in geometry rather than albedo. The presence of texture violates this constraint and it is assumed that the difficulties in the recovery of object shape in the shading plus texture case result from problems in separating changes in brightness due to geometry from those due to surface markings.

CONCLUSION

Overall, the introduction and enhancement of binocular disparity cues improved slant discrimination. Curvature discrimination is best in the case of shaded stereoscopic displays. There is no evidence for the kind of sequential dependencies predicted by the symbolic pipeline architecture of shape processing and the data points to independent encoding systems for surface situation and surface shape.

REFERENCES

- Barrow, H. G. & Tenenbaum, J. M. (1978). Recovering intrinsic scene characteristics from images. In Hanson, A. & Riseman, E. M. (Eds), *Computer vision systems* (pp. 3-26). New York: Academic Press.
- Blake, A., Zisserman, A. & Knowles, G. (1985). Surface descriptions from stereo and shading. *Image and Vision Computing*, 3, 183-191.
- Brady, M., Ponce, J., Yuille, A. & Asada, H. (1985). Describing surfaces. In Hanafusa, H. & Inove, H. (Eds), *Robotics research* (pp. 5-16). Cambridge, Mass.: MIT Press.
- Braunstein, M. L. (1976). *Depth perception through motion*. New York: Academic Press.
- 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.
- Bülthoff, H. H. & Mallot, H. A. (1990). Integration of stereo, shading and texture. In Blake, A. & Troscianko, T. (Eds), *AI and the eye* (pp. 119-146). Chichester: Wiley.
- Cagenello, R. & Rogers, B. J. (1993). Anisotropies in the perception of stereoscopic surfaces: The role of orientation disparity. *Vision Research*, 33, 2189-2202.
- Carman, G. J. & Welch, L. (1992). Three-dimensional illusory contours and surfaces. *Nature*, 360, 585-587.
- Erens, R. G. F., Kappers, A. M. L. & Koenderink, J. J. (1993). Perception of local shape from shading. *Perception & Psychophysics*, 54, 145-156.
- Foley, J. D., van Dam, A., Feiner, S. K. & Hughes, J. F. (1990). *Computer graphics* (2nd edn). Reading, Mass.: Addison Wesley.
- Frisby, J. P. & Pollard, S. B. (1991). Computational issues in solving the stereo correspondence problem. In Landy, M. & Movshon, J.A.

- (Eds), *Computational models of visual processing* (pp. 331-356). Cambridge, Mass.: MIT Press.
- Gibson, J. J. (1950). *The perception of the visual world*. Boston: Houghton Mifflin.
- Horn, B. K. P. (1975). Obtaining shape from shading information. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 115-155). N.Y.: McGraw-Hill.
- Johnston, A. (1992). Object constancy in face processing: Intermediate representations and object forms. *Irish Journal of Psychology*, 13, 425-438.
- Johnston, A. & Passmore, P. J. (1994). Shape from shading I: Surface curvature and orientation. *Perception*, 23, 169-189.
- Johnston, A., Passmore, P. J. & Morgan, M. J. (1991). Curvature discrimination thresholds for shaded and textured surfaces. *Investigative Ophthalmology and Visual Science*, 32, 1179.
- Koenderink, J. J. (1990). *Solid shape*. Cambridge, Mass.: MIT Press.
- Koenderink, J. J. & van Doorn, A. J. (1980). Photometric invariants related to solid shape. *Optica Acta*, 27, 981-996.
- Koenderink, J. J., van Doorn, A. J. & Kappers, A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, 52, 487-496.
- Marr, D. (1978). Representing visual information—a computational approach. In Hanson, A. & Riseman, E. M. (Eds), *Computer vision systems* (pp. 61-80). New York: Academic Press.
- Marr, D. (1982). *Vision*. San Francisco, Calif.: Freeman.
- Mingolla, E. & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics*, 53, 137-151.
- Perrone, J. A. (1980). Slant underestimation: A model based on the size of the viewing aperture. *Perception*, 9, 285-302.
- Phong, B.-T. (1975). Illumination for computer generated images. *Communications of the ACM*, 18, 311-317.
- Reichel, F. D. & Todd, J. T. (1991). Metric knowledge of three dimensional surface structure. *Investigative Ophthalmology and Visual Science*, 32, 1179.
- Rogers, B. & Cagenello, R. (1989). Disparity curvature and the perception of three-dimensional surfaces. *Nature*, 339, 135-137.
- Stevens, K. A. (1983). Slant-tilt: The visual encoding of surface orientation. *Biological Cybernetics*, 46, 183-195.
- Stevens, K. (1984). On gradients and texture 'gradients'. *Journal Experimental Psychology: General*, 113, 217-220.
- Todd, J. T. & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, 4, 643-657.
- Watt, R. J. & Andrews, D. P. (1981). APE: Adaptive Probit Estimation of psychometric functions. *Current Psychological Reviews*, 1, 205-214.

Acknowledgement—This work was supported by a grant from the SERC Image Interpretation Initiative.