Perception, 1994, volume 23, pages 191 – 200

Shape from shading. II. Geodesic bisection and alignment

Alan Johnston, Peter J Passmore
Department of Psychology, University College London, Gower Street, London WC1E 6BT, UK
Received 16 September 1992, in revised form 22 July 1993

Abstract. Pattern-acuity tasks have provided valuable information about the precision with which the visual system can make judgments about relative spatial position in two-dimensional images. However, outside the laboratory the visual system is habitually faced with the more difficult task of making positional judgments within a three-dimensional spatial environment. Thus our perceptual systems for representing surface shape also need to support the recovery of the location and disposition of features in a three-dimensional space. An investigation of the precision of three-dimensional position judgments in two spatial-judgment tasks, arc length bisection along geodesics and geodesic alignment, is reported. The spatial-judgment tasks were defined with reference to a sphere rendered by means of ray-casting techniques. The presence of shading and texture cues had no effect on discrimination thresholds in either task. Observers' constant errors were generally less than the just noticeable distance, demonstrating that the observers can perform these positional judgment tasks without substantial bias. It is argued that there is no explicit computation of arc length on the basis of shading and texture information and that surface-orientation information cannot be used as a reference in geodesic-alignment tasks. The results raise questions about the utility of a representation of surface orientation in the human visual system.

1 Introduction

In part I (this issue) we considered a number of schemes for the representation of surface geometry and argued that the shape parameter of Koenderink's local-shape index (Koenderink 1990) provides the most useful basis for an intermediate representation of visible surfaces because it provides invariance over transformations of scale and orientation. These invariant properties provide for a degree of stability in the description of object shape under transformation which is unavailable in a representational system like Marr and Nishihara's 2.5-D sketch (Marr and Nishihara 1978) which describes surface range and orientation.

Although a shape-index description would provide stability in object perception, the visual system would still need to detect and represent the orientation of objects relative to some reference frame. Experiments on surface orientation from shading and texture cues have not provided much support for a local representation of surface orientation like that envisaged by Marr and Nishihara (1978), Barrow and Tenenbaum (1978), and Horn (1975). Subjects find great difficulty in reporting the slant and tilt components of surface orientation (Mingolla and Todd 1986) and observers' ability to detect a change in higher-order properties of a surface, such as surface curvature, is much better than would be expected from their orientation discrimination thresholds (see part I, this issue).

The interest in representation in cognitive science comes from a realisation that the form that information takes can radically affect the ease with which procedures utilising that information can be implemented. Therefore it is natural to evaluate a scheme for the representation of surface geometry in terms of its utility as well as its adequacy. One advantage of a representation of the kind envisaged by Marr is that it could be utilised to make explicit other aspects of surface geometry. Presumably, we need to represent three-dimensional (3-D) information in a form that will support judgments about the location of objects or object features in a three-dimensional space as well as supporting shape perception. The tangent vectors to the surface allow the specification of the first fundamental form (O'Neill 1966). In differential geometry the first fundamental form, which describes the metric of a surface patch, is the usual route to the computation of properties which are intrinsic to the surface, like distances, angles, surface areas, and Gaussian curvature. It is possible that tasks which draw on subjects' perceptual skills in this area would provide evidence for a dense map of surface orientation in the human visual system.

There are a number of ways in which the pictorial information could conceivably be used to support judgments about spatial intervals along surfaces. A straightforward approach would be to use the information in the interval to compute its extent in three-dimensional space. For shading or texture we could first recover information about surface gradients. Arc length can then be computed by using the first fundamental form of the surface (O'Neill 1966). The first fundamental form for a Monge patch, in which a surface $\boldsymbol{x}(x,y)$ is parameterised by the image coordinates, is given by

$$I = \mathrm{d}s^2 = E\,\mathrm{d}x^2 + 2F\,\mathrm{d}x\mathrm{d}y + G\,\mathrm{d}y^2$$

where

$$E = \boldsymbol{x}_x \cdot \boldsymbol{x}_x\,, \quad F = \boldsymbol{x}_x \cdot \boldsymbol{x}_y\,, \quad E = \boldsymbol{x}_y \cdot \boldsymbol{x}_y\,.$$

The arc length of the image of the curve $\boldsymbol{x} = \boldsymbol{x}(x(t),\, y(t))$, $a \leqslant t \leqslant b$, can be computed from

$$s = \int_a^b \left[E\left(\frac{\mathrm{d}x}{\mathrm{d}t}\right)^2 + 2F\frac{\mathrm{d}x}{\mathrm{d}t}\frac{\mathrm{d}y}{\mathrm{d}t} + G\left(\frac{\mathrm{d}y}{\mathrm{d}t}\right)^2 \right]^{1/2} \mathrm{d}t\,.$$

Texture could be utilised in a more direct way. One of the most robust findings in the pattern-acuity literature is that discrimination thresholds in an interval-discrimination task increase in proportion to the size of the interval: the Weber relation for spatial-interval discrimination (Andrews and Miller 1978; Burbeck 1987; Levi and Klein 1983, 1989; Morgan and Watt 1989; Westheimer 1979). If the pattern within an interval can be used to supply a local reference then the task can be redefined and performance can improve. This is essentially the property we exploit when we use a rule to bisect an interval. We align the divisions on the rule with the endpoints of the interval and make the mark by using local spatial judgments. As Gibson (1979, page 162) pointed out, there are "equal amounts of texture for equal amounts of terrain", allowing surface texture to act as a measure marking off equal divisions along the surface.

Discrimination thresholds are a measure of an observer's uncertainty in performing a discrimination task. It is assumed that an explicit computation of arc length would reduce uncertainty and the variability of subjects' decisions when compared with image-based strategies for performing the task. We would expect performance on surface-interval judgments should improve in the presence of depth cues when compared with a control condition in which only minimal information about the surface orientation is available. In the first experiment we investigated the effect of addition and removal of depth cues in an geodesic-bisection task defined on shaded and textured computer-generated spheres.

2 General methods

2.1 *Subjects*

The two authors, AJ and PP, served as subjects in most of the experiments. Both had extensive practice on these tasks. Subjects AG and AF were not aware of the purpose of the study. All subjects had normal or well-corrected vision.

2.2 *Stimulus generation and display*
Full details of stimulus generation and display are given in part I (this issue). An image of a sphere was constructed by ray casting (Foley et al 1990). The sphere was rendered by means of a Phong illumination model,

$$P = sI_a + sI_p(\boldsymbol{N} \cdot \boldsymbol{L}) + gI_p(\boldsymbol{H} \cdot \boldsymbol{N})^n,$$

where P is the computed brightness, s is the albedo, I_a is the intensity of ambient illumination, I_p is the intensity of direct illumination, and g controls the proportion of light reflected specularly. $\boldsymbol{N}$ and $\boldsymbol{L}$ are the surface-normal and light-source-direction unit vectors, respectively, and $\boldsymbol{H}$ is the unit vector which bisects $\boldsymbol{L}$ and the line of sight. The spread of specular reflection is controlled by the parameter n. Texture could be added to the sphere by a texture-mapping technique. To generate the texture map we filtered a 256×256 pixel random grey-level image using an ideal filter which had a 1 octave bandwidth and a centre frequency of 28.8 cycles per image. The stimuli were displayed on a 19-inch Sony Trinitron monitor screen under the control of a SUN Sparcstation 330. In order to linearise the display a lookup table of luminance values was determined with a microphotometer and was used to control stimulus brightness. In some experiments the shaded and textured spheres were replaced by a circular field of uniform brightness (42 cd m^{-2}).

2.3 *Procedure*
Three small spherical bumps were attached to the sphere (figure 1). The centres of the bumps were coincident with the surface of the sphere. Only those parts of the bumps which lay outside the main body of the sphere were visible. The bumps were located with subpixel accuracy because bump position was assigned geometrically and the brightness values in the region were computed by ray casting. In the specification of bump position the sphere is considered to be oriented with the north pole at the upper margin of the image. The two outer bumps were positioned either on the equator or along meridians of longitude. For bumps positioned along the equator the subjects' task was to decide whether the central bump bisected the arc length of the geodesic connecting the two outer bumps. This task is referred to as geodesic bisection. For bumps positioned along meridians the subjects' task was to decide whether a more central bump lay to the left or right of the geodesic connecting the two outer bumps. This task is referred to as geodesic alignment.

Thresholds were measured with an adaptive method of constant stimuli, APE (Watt and Andrews 1981). Discrimination threshold is defined as the standard deviation of the error distribution and corresponds to the 84% point on the psychometric function.

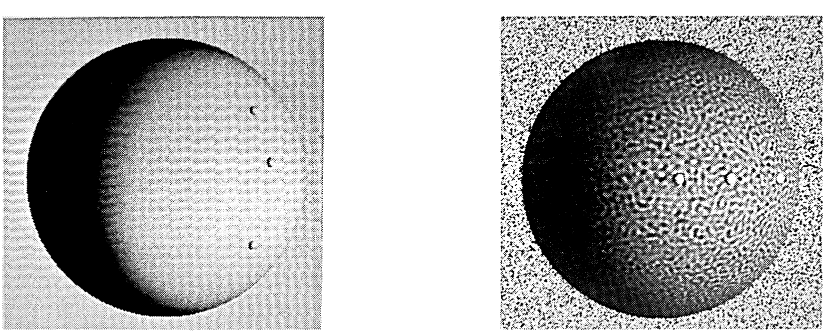

(a) (b)

Figure 1. (a) An example of one of the stimuli from the geodesic-alignment task in a smooth-shading condition. (b) One of the stimulus conditions in the arc-length-bisection task in which texture has been added to the sphere.

Thresholds were based on 64 individual trials. Each data point is the root mean square of at least four separate threshold determinations and is therefore based on over 256 trials. The standard deviations of the individual threshold determinations provide a measure of dispersion. Stimuli were viewed monocularly with the dominant eye. A chin rest was used to help stabilise the head and a viewing tube reduced the field of view to the central part of the display containing the rendered sphere. In the geodesic-bisection task the sphere was displayed on a random grey-level background. In the alignment tasks we used a homogeneous grey field.

3 Experiment 1: Geodesic bisection

In the first experiment we measured thresholds for arc length bisection, an analogue to the familiar two-dimensional (2-D) interval bisection task. A computer-generated sphere with a diameter of 7.5 cm was viewed from a distance of 100 cm. The sphere was drawn by using perspective projection with rays cast from the stationpoint. Lighting parameters were chosen to model the effects of point-source illumination on a perfect diffuse reflector or Lambertian surface. The illumination parameters were $I_a = 0.1$, $I_p = 0.7$, $s = 1$, $g = 0$. The point source was positioned above {0, 100, 100 cm} the observer. In this task three bumps were positioned along the equator. The interval between outer bumps subtended an angle of 30 deg at the origin of the sphere. When symmetrically placed around the centre of the display, the interval subtended an angle of 1.15 deg at the eye. Each bump subtended 10 min. Subjects were asked to indicate whether the central bump was to the right or left of the midpoint of a geodesic arc connecting the two outer bumps.

Discrimination thresholds for arc length bisection were measured as a function of meridional position in three depth-cue conditions. The bumps were drawn either on a smoothly shaded sphere, on a shaded and textured sphere, or on a disc of uniform brightness. In the control condition observers were asked to imagine the bumps were on the surface of a sphere. Trials for the four meridians in each depth-cue condition were interleaved. Note that the band-limited texture has the effect of marking out units of arc length along the surface. The location of the three bump configuration was perturbed by a factor of ±1 deg (sphere units), so that subjects could not use the relative position of the bumps with respect to the texture as a cue. The jitter range is around 8 times threshold for the 0°-meridian condition.

If a representation of surface orientation can be utilised in the computation of arc length, subjects can simply compare the lengths for the two intervals in order to arrive at a decision. If there is no explicit computation of arc length, subjects would have to decide from trial to trial whether the angular subtenses of the two intervals with respect to the eye were in the correct ratio for that position on the sphere. We would expect that the latter strategy would lead to greater variability in subjects' responses and higher thresholds. Since a dense map of surface-orientation information would not be available in the control condition we would expect thresholds to be higher than in the conditions with shading and texture cues.

Figure 2 shows the results for two subjects. Thresholds were lower for shading at the 45° meridian for PP and at the 0° meridian for AJ but we found no systematic effects of the presence of 3-D cues on the interval-bisection task across both subjects. Thresholds were fairly constant as a function of meridional angle. Constant errors tended to be less than the just noticeable difference but varied between observers and indeed between eyes for a given individual (Lindblom and Westheimer 1989), suggesting that the underlying cause of constant errors in bisection tasks may be geometric aberration in the eye. There were no substantial effects of conditions on constant errors.

Observers are able to bisect an arc on a sphere defined by shading and texture with a considerable degree of accuracy. However, for some conditions, observers are equally precise in the control condition, in which the only visual indication of shape provided was a bounding contour. The presence of shading and texture did not lead to the reduction in discrimination threshold we might have expected if subjects were able to compute arc length on the basis of these cues or if texture could be used as a guide to the extent of an interval (Craven and Watt 1989; Gibson 1979).

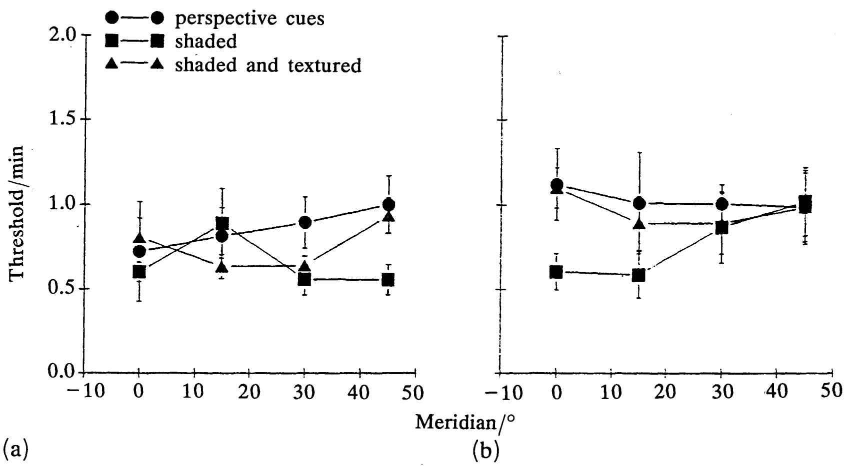

Figure 2. Discrimination thresholds for the arc length bisection at a range of meridians of longitude in three depth-cue conditions: perspective cues, no shading or texture; shading; and shading and texture. (a) Subject PP; (b) subject AJ. In this and subsequent figures vertical bars represent ±1 standard deviation.

4 Experiment 2: Effect of lighting direction on geodesic alignment

The finding that the presence of shading has little effect on arc length discrimination is somewhat surprising. If surfaces are described by a map of surface gradients or normals, as in Marr and Nishihara's 2.5-D sketch (Marr and Nishihara 1978), then it would appear that surface gradients cannot be used to compute arc length. However, a dense representation of surface orientation may be utilised in other spatial tasks.

In 2-D spatial vision, investigations of collinearity judgments by means of dots and lines have a long history. The most natural extension of this paradigm to solid-shape perception involves judgments of coplanarity. Along geodesics on the sphere, surface normals are coplanar. Formally, vectors are coplanar if the triple scalar vector product equals zero. It is not clear whether a biological mechanism could compute the triple product; however, along the meridians of longitude on the sphere the first partial derivative of z with respect to x, the p parameter of gradient space (Horn 1977; Huffman 1971), is constant. Thus a representation of surface orientation in terms of surface gradients could support geodesic-alignment judgments along meridians of longitude. For example, if the meridian of interest is defined by two dots placed symmetrically above the equator, the location of a third dot will be on one side of the meridian if the gradient in the x direction is greater at that point than at the meridian, and on the opposite side if the gradient at that point is less than at the meridian. The logic of the experiment is that a surface-gradient representation could, in principle, provide an explicit local reference for the positional-judgment task, resulting in a reduction in discrimination threshold compared with that in a control condition in which strong 3-D cues to shape are eliminated.

In experiment 2 we investigated the effects of meridional position and the direction of illumination on thresholds for geodesic alignment. We attempted to manipulate the salience of the shading cue by varying the direction of illumination. It is generally assumed that the visual system interprets scenes as being illuminated from above, although Reichel and Todd (1990) have demonstrated that surfaces show reversals with stimulus inversion for images in which the simulated light source was positioned at the point of observation.

As in experiment 1 we generated a shaded sphere which was viewed at a distance of 100 cm. Lighting parameters were chosen to model the effects of point-source illumination on a perfect diffuse reflector or Lambertian surface. The illumination-direction vector was always at an angle of 45° to the line of sight. Thresholds were measured for bumps placed on five meridians of longitude (0°, 15°, 30°, 45°, and 60°). The bumps had a diameter of 10 min. The outer two bumps were positioned symmetrically about the equator on the 15° parallels of latitude. The central dot was positioned on the equator. To change the direction of illumination the full display was rotated by means of a dove prism. For each light-source position, thresholds at all meridians were measured in one sitting, with individual trials interleaved.

The results are shown in figure 3. There is a tendency, more evident in the data for PP, for thresholds to increase slightly with meridional angle up to 45°. There are no systematic effects of lighting direction on alignment thresholds. This manipulation may only have marginal effects on the 3-D representation of the surface and therefore we decided to repeat the experiment with a range of backdrops. As in experiment 1, subjects may not be utilising a map of surface normals to perform the task, in which case adding or removing 3-D cues should have no effect.

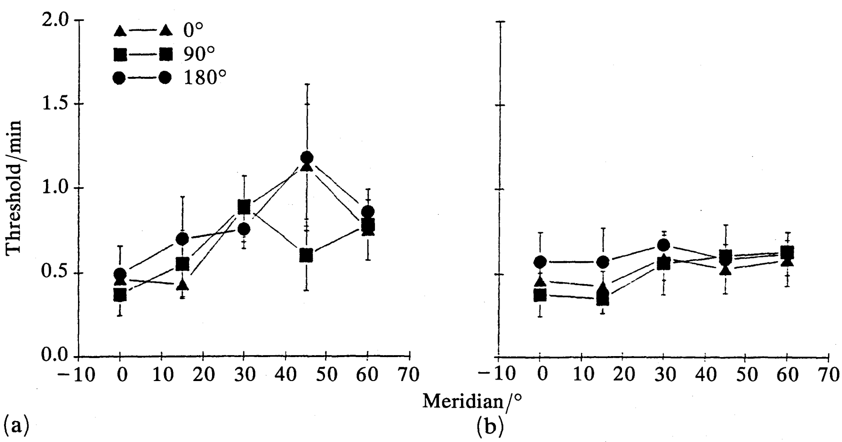

Figure 3. Discrimination thresholds for geodesic alignment at a range of meridians of longitude in three depth-cue conditions: 0° tilt, illuminated from above; 90° tilt, illuminated from the right; and 180° tilt, illuminated from below. (a) Subject PP; (b) subject AJ.

5 Experiment 3: Effect of depth cues on geodesic alignment

The methods used in experiment 3 were identical to those used in experiment 2. Geodesic-alignment thresholds were measured for bumps presented on a uniform disc. The size of the bumps either varied with position in the depth plane as in experiment 2 (the perspective-cues condition), or were constant in size (the no-perspective-cues condition). In two other conditions texture cues were added to the shaded sphere and in one of those conditions the lighting model included a specular component ($g = 0.37$, $n = 32$). In figure 4a thresholds for the perspective-cues and the no-perspective-cues condition are plotted, together with the data from experiment 2. Thresholds for all three lighting conditions in experiment 2 have been combined. Clearly, replacing the shading object with a uniform disc has no effect on discrimination threshold. Adding texture and a specular highlight appeared to produce the impression of a more solid object but had no substantial effect on thresholds (figure 4b). Data for AJ (not shown) confirmed the finding that there was no effect of the addition or removal of depth cues. Experiments 2 and 3 do not support the proposal that surface-gradient information can improve performance in the geodesic-alignment task. If subjects were able to derive a local reference from surface-gradient information, removal of shading cues should have degraded performance.

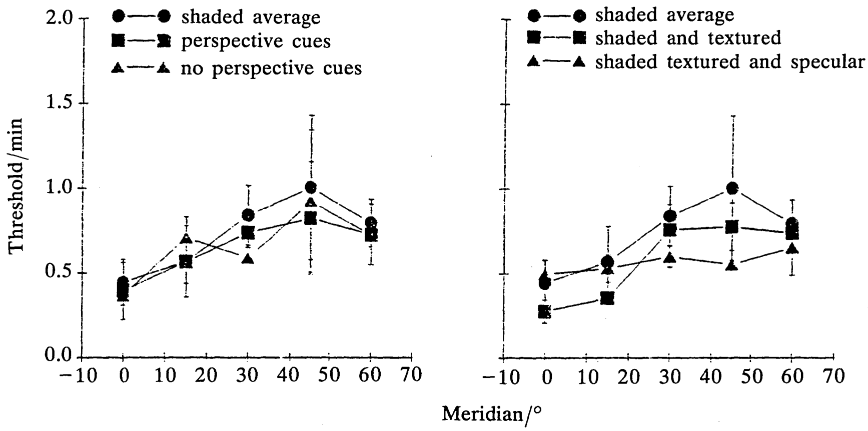

(a) (b)

Figure 4. (a) Discrimination thresholds for geodesic alignment at a range of meridians of longitude in three depth-cue conditions. The circles are the average shaded-sphere thresholds from experiment 2, the squares are the data for perspective cues only, and the triangles are the data for the no-depth-cues condition. (b) Discrimination thresholds for geodesic alignment with additional depth information. The circles are the average shaded-sphere thresholds replotted from experiment 2, the squares are the data for shading and texture cues, and the triangles are data for shaded, textured spheres with a specular highlight. Subject PP.

6 Experiment 4: Effect of longitudinal jitter on geodesic alignment

In the depth-cue conditions subjects could perform the task by aligning the bumps with reference to the level sets of the gradient-space parameter, p, the component of the surface gradient in the x direction. These level sets essentially mark out the meridians of longitude on the image. It would appear that subjects used a strategy based on a 2-D image rather than using any surface-orientation information that might be available. However, we might be able to bias subjects towards a strategy based on surface gradients by adding stimulus-related noise which could affect a strategy based on a 2-D image but which would not affect a surface-based strategy. In experiment 4 the position of the central bump along the meridian of longitude was randomised from trial to trial within the range $\pm 15°$. The separation of the outer bumps was increased to $\pm 30°$ about the equator. Note that the value of the gradient at the points indicated by the bumps is unaffected by this manipulation. The sphere was illuminated from the right {100, 0, 100 cm}.

Figure 5 shows three bump-alignment thresholds as a function of meridional position. Each data point shows the average thresholds for three subjects. Thresholds are higher overall in this experiment because of the increased separation of the outer dots. With longitudinal jitter there is a marked increase in alignment thresholds with meridional position. Since longitudinal jitter should not compromise a strategy based on a comparison of the magnitudes of the surface gradients at the positions marked by the bumps it is unlikely that subjects were able to use this procedure. Again there was no effect of the presence of the shading cue.

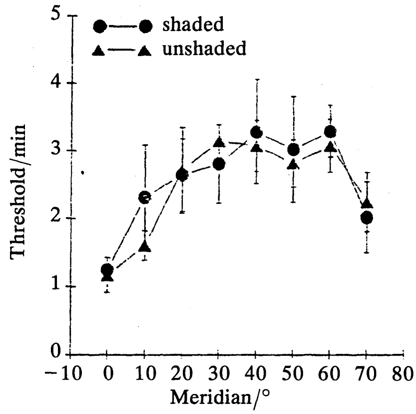

Figure 5. Discrimination thresholds for geodesic alignment at a range of meridians of longitude. The latitude of the central bump was perturbed within the range ±15°. Each data point shows the averaged thresholds for three subjects in two conditions: with shading and without shading.

7 Discussion

7.1 *Position judgments on computer-generated surfaces*

Watt (1991) points out that there is no clear or consensus view of how the distance between two points is measured by the human visual system. Craven and Watt (1989) suggest that the visual system might count the number of zero-bounded responses of a range of filters at various scales. Methods of computing the length of an interval which are based on the content of the interval have the advantage that they are readily generalisable to 3-D surfaces. However, we found no systematic difference in discrimination thresholds for empty and textured intervals. Gibson's (1950, 1979) assertion that texture provides a means of perceiving distance along surfaces is not supported and there is no indication that we can compute surface distance from surface gradients.

In both the geodesic-bisection and the geodesic-alignment task constant errors were generally less than the discrimination threshold, demonstrating that observers followed instructions and could veridically judge the null-cue locations without substantial bias, but there is little evidence for the explicit computation of arc length, collinearity, or coplanarity. Bisection is not required for optimal performance in spatial-interval-judgment tasks. Observers can divide an internal interval into an appropriate ratio with a precision which is equal to that for bisection (DeValois et al 1990). In the geodesic-alignment task (experiments 2 and 3), discrimination thresholds increased slightly (PP) or marginally (AJ) for eccentric meridians, where the bumps lay on an arc connecting the poles and the outer dots, when compared with the 0°-meridian conditions, in which the bumps fell on a straight line in the image plane.

Acuity for spatial interval and alignment for surfaces defined by shading and texture appears to be limited by the same factors that limit 2-D spatial judgments. The main factor affecting acuity in bisection tasks and alignment tasks is the separation of elements of the display (Andrews and Miller 1978; Burbeck 1987; Levi and Klein 1983, 1989; Morgan and Watt 1989). Increases in discrimination threshold with separation has been attributed to the observation that parts of the display must lie in the periphery of the visual field for large separations (Levi et al 1988). The grain of spatial analysis may be a contributing factor to the Weber law for spatial-interval discrimination but Morgan and Watt (1989) have shown that discrimination thresholds for interval judgments increase when the target stimuli are located along isoeccentricity contours in the visual field where spatial grain can be considered to be invariant.

An observer's bias in interval-judgment tasks can vary without any appreciable effect on discrimination thresholds (Morgan et al 1990; Westheimer 1979, 1981).

It is likely that in the surface-position tasks subjects use whatever information is available in the display to make an intellectual judgment about the required positions of the relevant intervals. Thesholds are then determined by the accuracy with which relative position can be encoded in the image.

7.2 *Do we have a dense map of surface orientation?*
We argued that the extent of an arc on a surface could be computed from surface-gradient information. It was assumed that if arc length was computed explicitly on the basis of a dense map of surface orientation then thresholds would be lower in the presence of shading and texture cues to surface shape. In addition we proposed that the level sets of the surface-gradient parameter, p, could provide a local reference in the geodesic-alignment task, giving improved performance in shaded-images relative to control conditions and indifference to longitudinal jitter. However, we found no evidence for the utilisation of these strategies in the spatial-acuity tasks.

There is little empirical evidence, from work on shading and texture cues, for the existence of a dense, precise map of surface orientation (Barrow and Tenenbaum 1978; Marr and Nishihara 1978) in the human visual system. Surface-orientation discrimination thresholds are at best around 3 deg in the conditions explored with our techniques (Johnston and Passmore, this issue). Orientation discrimination thresholds are too high to explain our abilities to discriminate changes in surface curvature (see part I, this issue) and there is little evidence for the use of surface-orientation information in surface-position judgments. Thresholds for arc length bisection, a task which could also be accomplished by bisecting the surface normals at the outer bump positions, were around 15 min (sphere units). It is difficult to imagine that this level of performance could be achieved on the basis of a symbolic representation of surface orientation, given orientation thresholds of around 10° for central positions on the sphere.

It is possible that a representation of surface orientation exists but is so imprecise as to be ineffective in these tasks. Alternatively, the representation of spatial information may be inflexible. The computational architecture of the visual system may not allow a map of surface orientation to be used in the service of 3-D spatial judgments. However, if this is the case, we have to question its utility. Marr (1978, 1982) considered that advantages of the 2.5-D sketch to be that (a) it provides an intermediate stage in representing the scene based on what it is possible to compute rather than what is desirable; (b) it makes surface orientation explicit; (c) it introduces mutual constraints on surface information, eg the orientation map can be derived from the height map by differentiation; and (d) it allows integration of separate depth modules. However, in the case of shape from shading and texture, we find that curvature discrimination does not appear to be based on the encoding of surface orientation, the encoding of surface orientation lacks precision, and there is little evidence of a surface-orientation map being put to further use, and therefore it is not clear that a surface-orientation description is the most appropriate point at which to combine depth information from different sources.

8 Conclusion

We examined the utility of a dense map of surface orientation as a representation of surface shape. We argued that a surface-orientation map could support precise judgment of distance along geodesics on the surface and alignment with respect to geodesics on the surface. However we found no consistent or substantial effects of removing texture or shading cues on performance in geodesic-bisection and geodesic alignment tasks. It would appear that subjects performed these tasks on the basis of strategies based on 2-D images. This has some practical implications for the design of visual displays. Increasing the fidelity of pictorial depth information is unlikely to improve observers' abilities to judge spatial position in a 3-D scene. The findings demonstrate that the human visual system, which has evolved piecemeal to perform certain tasks, cannot be considered an infinitely flexible, ideal visual computer. There are constraints on the precision with which abstract properties of the world can be represented and limits on the utilisation of available information.

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

References

Andrews D P, Miller D T, 1978 "Acuity for spatial separation as a function of stimulus size" *Vision Research* **18** 615-619

Barrow H G, Tenenbaum J M, 1978 "Recovering intrinsic scene characteristics from images", in *Computer Vision Systems* Eds A Hanson, E M Riseman (New York: Academic Press) pp 3-26

Burbeck C A, 1987 "Position and spatial frequency in large-scale localisation judgments" *Vision Research* **27** 417-427

Craven B, Watt R, 1989 "The use of fractal image statistics in the estimation of lateral spatial extent" *Spatial Vision* **4** 223-239

DeValois K K, Lakshminarayanan V, Nygaard R, Schlussel S, Slandky J, 1990 "Discrimination of relative spatial position" *Vision Research* **30** 1649-1660

Foley J D, Dam A van, Feiner S K, Hughes J F, 1990 *Computer Graphics* 2nd edition (Reading, MA: Addison-Wesley)

Gibson J J, 1950 *The Perception of the Visual World* (Boston, MA: Houghton Mifflin)

Gibson J J, 1979 *The Ecological Approach to Visual Perception* (Boston, MA: Houghton Mifflin)

Horn B K P, 1975 "Obtaining shape from shading information", in *The Psychology of Computer Vision* Ed. P H Winston (New York: McGraw-Hill) pp 115-155

Horn B K P, 1977 "Understanding image intensities" *Artificial Intelligence* **8** 201-231

Huffman D A, 1971 "Impossible objects as nonsense sentences" *Machine Intelligence* **6** 295-323

Koenderink J J, 1990 *Solid Shape* (Cambridge, MA: MIT Press)

Levi D M, Klein S A, 1983 "Spatial localization in normal and amblyopic vision" *Vision Research* **23** 1005-1017

Levi D M, Klein S A, 1989 "Both separation and eccentricity can limit precise position judgements: a reply to Morgan and Watt" *Vision Research* **29** 1463-1469

Levi D M, Klein S A, Yap Y L, 1988 "'Weber's Law' for position: unconfounding the role of separation and eccentricity" *Vision Research* **28** 597-603

Lindblom B, Westheimer G, 1989 "Binocular summation of hyperacuity tasks" *Journal of the Optical Society of America A* **6** 585-589

Marr D, 1978 "Representing visual information—a computational approach", in *Computer Vision Systems* Eds A Hanson, E M Riseman (New York: Academic Press) pp 61-80

Marr D, 1982 *Vision* (San Francisco, CA: W H Freeman)

Marr D, Nishihara H K, 1978 "Representation and recognition of the spatial organisation of three-dimensional shape" *Proceedings of the Royal Society of London, Series B* **200** 269-294

Mingolla E, Todd J T, 1986 "Perception of solid shape from shading" *Biological Cybernetics* **53** 137-151

Morgan M J, Watt R J, 1989 "The Weber's Law for distance is not an artifact of eccentricity" *Vision Research* **29** 1457-1462

Morgan M J, Hole G J, Glennerster A, 1990 "Biases and sensitivities in geometrical illusions" *Vision Research* 30 1793-1810

O'Neill B, 1966 *Elementary Differential Geometry* (New York: Academic Press)

Reichel F D, Todd J T, 1990 "Perceived depth inversion of smoothly curved surfaces due to image orientation" *Journal of Experimental Psychology: Human Perception and Performance* **16** 653-664

Watt R, 1991 *Understanding Vision* (London: Academic Press)

Watt R J, Andrews D P, 1981 "APE: Adaptive Probit Estimation of psychometric functions" *Current Psychological Review* **1** 205-214

Westheimer G, 1979 "The spatial sense of the eye" *Investigative Ophthalmology and Visual Science* **17** 893-912

Westheimer G, 1981 "Visual hyperacuity", in *Progress in Sensory Physiology* Volume 1, Eds P Gouras, E R Kandel, M Klein, H W Kosterlitz, A T McKnight, E Shapiro, G Westheimer, E Zrenner (New York: Springer) pp 1-30